21st International ACM Symposium on High-Performance Parallel and Distributed Computing

HPDC'12

Delft, the Netherlands

18-22 June 2012

Program

General Chair:
Dick Epema, Delft University of Technology, Delft, the Netherlands

Program Co-chairs:
Thilo Kielmann, Vrije Universiteit, Amsterdam, the Netherlands
Matei Ripeanu, The University of British Columbia, Vancouver, Canada

Posters Chair:
Ana Varbanescu, Delft University of Technology, Delft, the Netherlands

Workshops and Industry Session Chair:
Alexandru Iosup, Delft University of Technology, Delft, the Netherlands

HPDC 2012 Workshop Overview

Monday June 18	
09:00-17:30 Room 1.160	ScienceCloud: 3rd Workshop on Scientific Cloud Computing
11:00-16:30 Room 1.170	MapReduce'12: The Third International Workshop on MapReduce and its Applications
08:30-17:30 Room 01.180	ECMLS2012: 3rd International Emerging Computational Methods for the Life Sciences Workshop

Tuesday June 19	
09:00-17:30 Room 1.160	Astro-HPC: Workshop on High-Performance Computing for Astronomy
09:00-17:00 Room 1.170	DIDC: Fifth International Workshop on Data-Intensive Distributed Computing
09:00-17:10 Room 01.180	VTDC-2012: 6th International Workshop on Virtualization Technologies in Distributed Computing

Schedule of breaks on Monday and Tuesday:

Morning break:	between 10:00-11:00 (depends on workshop)
Lunch:	12:30-13:30
Afternoon break:	between 15:00 and 15:30 (depends on workshop)

All breaks and lunches of the workshops and the main conference, and the poster session and the reception of the main conference are in room 0.030 on the ground floor.

All sessions of the main conference are in room 0.160 on the ground floor.

HPDC 2012 Conference Program Overview

Wednesday June 20	
08:45 – 09:00	Conference Opening
09:00 - 10:00	Session 1: Keynote 1
10:00 – 10:20	Break
10:20 – 12:00	Session 2: Virtualization
12:00 – 13:15	Lunch
13:15 – 14:55	Session 3: I/O
14:55 – 15:15	Break
15:15 – 16:05	Session 4: Industry Session
16:05 – 16:15	Short break
16:15 – 17:05	Session 5: GPUs
17:05 – 17:50	Session 6: Poster Presentations
18:00 – 19:00	Session 7: Posters + Reception

Thursday June 21	
09:00 – 10:00	Session 8: Keynote 2
10:00 – 10:20	Break
10:20 – 12:00	Session 9: Applications and Resources
12:00 – 13:15	Lunch
13:15 – 14:55	Session 10: MapReduce
14:55 – 15:15	Break
15:15 – 16:05	Session 11: Energy
16:05 – 16:15	Short break
16:15 – 17:15	Session 12: Panel on Energy Efficiency
18:00 – 19:00	Social Event (1): Visit to Museum De Prinsenhof
19:00 – 22:30	Social Event (2): Conference Dinner in De Prinsenkelder

Friday June 22	
09:00 – 10:00	Session 13: Achievement Award Talk
10:00 – 10:20	Break
10:20 – 11:35	Session 14: Networked Systems
11:35 – 11:45	Short break
11:45 – 12:45	Session 15: Panel on the value of HPDC
12:45	Conference Closing (with Best Paper Award)

Wednesday, June 20

08:45-09:00 **Conference Opening**

09:00-10:00 **Session 1: Keynote 1**
Chair: *Matei Ripeanu*

Putting "Big-data" to Good Use: Building Kinect
Mihai Budiu, Microsoft Research, Mountain View, USA

10:00-10:20 **Break**

10:20-12:00 **Session 2: Virtualization**
Chair: *Beth Plale*

vSlicer: Latency-aware Virtual Machine Scheduling via Differentiated-frequency CPU Slicing
Cong Xu (Purdue University), Sahan Gamage (Purdue University), Pawan N. Rao (Purdue University), Ardalan Kangarlou (NetApp), Ramana Kompella (Purdue University), Dongyan Xu (Purdue University)

Singleton: System-wide Page Deduplication in Virtual Environments
Prateek Sharma, Purushottam Kulkarni (IIT Bombay)

Locality-aware Dynamic VM Reconfiguration on MapReduce Clouds
Jongse Park, Daewoo Lee, Bokyeong Kim, Jaehyuk Huh, Seungryoul Maeng (KAIST)

Achieving Application-Centric Performance Targets via Consolidation on Multicores: Myth or Reality?
Lydia Y. Chen Chen (IBM Research Zurich Lab), Danilo Ansaloni (University of Lugano), Evgenia Smirni (College of William and Mary), Akira Yokokawa (University of Lugano), Walter Binder (University of Lugano)

12:00-13:15 **Lunch**

13:15-14:55 **Session 3: I/O**
Chair: *Dongyan Xu*

Enabling Event Tracing at Leadership-Class Scale through I/O Forwarding Middleware
Thomas Ilsche (Technische Universität Dresden), Joseph Schuchart (Technische Universität Dresden), Jason Cope (Argonne National Laboratory), Dries Kimpe (Argonne National Laboratory), Terry Jones (Oak Ridge National Laboratory), Andreas Knöpfer (Technische Universität Dresden), Kamil Iskra (Argonne National Laboratory), Robert Ross (Argonne National Laboratory), Wolfgang E. Nagel (Technische Universität Dresden), Stephen Poole (Oak Ridge National Laboratory)

ISOBAR Hybrid Compression-I/O Interleaving for Large-scale Parallel I/O Optimization
Eric R. Schendel (North Carolina State University), Saurabh V. Pendse (North Carolina State University), John Jenkins (North Carolina State University), David A. Boyuka (North Carolina State University), Zhenhuan Gong (North Carolina State University), Sriram Lakshminarasimhan (North Carolina State University), Qing Liu (Oak Ridge National Laboratory), Scott Klasky (Oak Ridge National Laboratory), Robert Ross (Argonne National Laboratory), Nagiza F. Samatova (North Carolina State University)

QBox: Guaranteeing I/O Performance on Black Box Storage Systems
Dimitris Skourtis, Shinpei Kato, Scott Brandt (University of California, Santa Cruz)

Towards Efficient Live Migration of I/O Intensive Workloads: A Transparent Storage Transfer Proposal
Bogdan Nicolae (INRIA), Franck Cappello (INRIA/UIUC)

14:55-15:15 Break

15:15-16:05 Session 4: Industry Session
Chair: *Alexandru Iosup*

Parallelization of C Programs through Dependency Analysis
Jos van Eijndhoven, CTO of Vector Fabrics, Eindhoven, the Netherlands

Parallel Programming for the Masses
Vincent Hendriksen, StreamComputing, Haarlem, the Netherlands

16:05-16:15 Short break

16:15-17:05 Session 5: GPUs
Chair: *Ana Varbanescu*

A Virtual Memory Based Runtime to Support Multi-tenancy in Clusters with GPUs
Michela Becchi (University of Missouri), Kittisak Sajjapongse (University of Missouri), Ian Graves (University of Missouri), Adam Procter (University of Missouri), Vignesh Ravi (Ohio State University), Srimat Chakradhar (NEC Laboratories America)

Interference-driven Scheduling and Resource Management for GPU-based Heterogeneous Clusters
Rajat Phull, Cheng-Hong Li, Kunal Rao, Hari Cadambi, Srimat Chakradhar (NEC Laboratories America)

17:05-17:50 Session 6: Poster Presentations
Chair: *Ana Varbanescu*

18:00-19:00 Session 7: Posters + Conference Reception

Thursday, June 21

09:00-10:00 **Session 8: Keynote 2**
Chair: *Thilo Kielmann*

Leveraging Renewable Energy in Data Centers: Present and Future
Ricardo Bianchini, Rutgers University, USA

10:00-10:20 **Break**

10:20-12:00 **Session 9: Applications and Resources**
Chair: *Alexandru Iosup*

Work Stealing and Persistence-based Load Balancers for Iterative Overdecomposed Applications
Jonathan Lifflander (UIUC), Sriram Krishnamoorthy (PNNL), Laxmikant V. Kale (UIUC)

Highly Scalable Graph Search for the Graph500 Benchmark
Koji Ueno (Tokyo Institute of Technology/JST CREST), Toyotaro Suzumura (Tokyo Institute of Technology/IBM Research Tokyo/JST CREST)

PonD : Dynamic Creation of HTC Pool on Demand Using a Decentralized Resource Discovery System
Kyungyong Lee (University of Florida), David Wolinsky (Yale University), Renato Figueiredo (University of Florida)

SpeQuloS: A QoS Service for BoT Applications Using Best Effort Distributed Computing Infrastructures
Simon Delamare (INRIA), Gilles Fedak (INRIA), Derrick Kondo (INRIA), Oleg Lodygensky (IN2P3)

12:00-13:15 **Lunch**

13:15-14:55 **Session 10: MapReduce**
Chair: *Carlos Maltzahn*

Understanding the Effects and Implications of Compute Node Related Failures in Hadoop
Florin Dinu, T. S. Eugene Ng (Rice University)

Optimizing MapReduce for GPUs with Effective Shared Memory Usage
Linchuan Chen, Gagan Agrawal (The Ohio State University)

CAM: A Topology Aware Minimum Cost Flow Based Resource Manager for MapReduce Applications in the Cloud
Min Li (Virginia Tech), Dinesh Subhraveti (IBM Almaden Research Center), Ali Butt (Virginia Tech), Aleksandr Khasymski (Virginia Tech), Prasenjit Sarkar (IBM Almaden Research Center)

Distributed Approximate Spectral Clustering for Large-Scale Datasets
Fei Gao (Simon Fraser University), Wael Abd-Almageed (University of Maryland), Mohamed Hefeeda (Qatar Computing Research Center)

14:55-15:15 **Break**

15:15-16:05 **Session 11: Energy**
Chair: *Henri Bal*

Exploring Cross-layer Power Management for PGAS Applications on the SCC Platform
Marc Gamell (Rutgers University), Ivan Rodero (Rutgers University), Manish Parashar (Rutgers University), Rajeev Muralidhar (Intel India)

Dynamic Adaptive Virtual Core Mapping to Improve Power, Energy, and Performance in Multi-socket Multicores
Chang Bae (Northwestern University), Lei Xia (Northwestern University), Peter Dinda (Northwestern University), John Lange (University of Pittsburgh)

16:05-16:15 **Short break**

16:15-17:15 **Session 12: Panel on Energy Efficiency**

Moderator: *Matei Ripeanu, The University of British Columbia, Vancouver, Canada*

Members: *Ricardo Bianchini (Rutgers University)*
Manish Parashar (Rutgers University)
Karsten Schwan (Georgia Tech)

18:00-19:00 **Social Event (1): Visit to Museum De Prinsenhof**

19:00-22:30 **Social Event (2): Conference Dinner in Restaurant De Prinsenkelder**

Friday, June 22

09:00-10:00 **Session 13: Achievement Award Talk**
Chair: *Dick Epema*

Reflections on 20 Years of Grid Computing
Ian Foster, University of Chicago and Argonne National Laboratory, USA

10:00-10:20 **Break**

10:20-11:35 **Session 14: Networked Systems**
Chair: *Renato Figueiredo*

VNET/P: Bridging the Cloud and High Performance Computing Through Fast Overlay Networking
Lei Xia (Northwestern University), Zheng Cui (University of New Mexico), John Lange (University of Pittsburgh), Yuan Tang (UESTC, China), Peter Dinda (Northwestern University), Patrick Bridges (University of New Mexico)

Massively-Parallel Stream Processing under QoS Constraints with Nephele
Björn Lohrmann, Daniel Warneke, Odej Kao (Technische Universität Berlin)

A Resiliency Model for High Performance Infrastructure Based on Logical Encapsulation
James Moore (The University of Southern California/EMC Corporation), Carl Kesselman (The University of Southern California)

11:35-11:45 **Short break**

11:45-12:45 **Session 15: Panel on the Value of HPDC**

Moderator: *Henri Bal, Vrije Universiteit, the Netherlands*

Members: *Renato Figueiredo (University of Florida)*
Ian Foster (University of Chicago and Argonne National Laboratory)
Rob van Nieuwpoort (Netherlands eScience Center)
Jon Weissman (University of Minnesota)

12:45-13:00 **Conference Closing (with Best Paper Award)**

Accepted Posters (based on full submitted papers)

Dynamic Binary Rewriting and Migration for Shared-ISA Asymmetric, Multicore Processors
Giorgis Georgakoudis (University of Thessaly), Dimitrios S. Nikolopoulos (Queen's University of Belfast)

Exploring the Performance and Mapping of HPC Applications to Platforms in the Cloud
Abhishek Gupta (UIUC), Laxmikant V. Kale (UIUC), Dejan S. Milojicic (HP labs, Palo Alto), Paolo Faraboschi (HP labs, Palo Alto), Richard Kaufmann (HP labs, Palo Alto), Verdi March (HP labs, Singapore), Filippo Gioachin (HP labs, Singapore), Chun Hui Suen (HP labs, Singapore), Bu-Sung Lee (HP labs, Singapore)

Fault Tolerant Data Intensive Algorithms
Mucahid Kutlu, Gagan Agrawal, Oguz Kurt (The Ohio State University)

P*: A Model of Pilot-Abstractions
Andre Luckow (CCT/LSU), Mark Santcroos (AMC/University of Amsterdam), Ole Weidner (CCT/LSU), Andre Merzky (CCT/LSU), Sharath Maddineni (CCT/LSU), Shantenu Jha (Rutgers University)

Performance Evaluation of Inter-thread Communication Mechanisms on Multicore/multithreaded Architectures
Davide Pasetto, Massimiliano Meneghin, Hubertus Franke, Fabrizio Petrini, Jimi Xenidis (IBM Research)

Coupling Task Progress for MapReduce Resource-Aware Scheduling
Jian Tan, Xiaoqiao Meng, Li Zhang (IBM T.J. Watson)

SMART-IO: SysteM-AwaRe Two-Level Data Organization for Efficient Scientific Analytics
Yuan Tian (Auburn University), Scott Klasky (Oak Ridge National Laboratories), Weikuan Yu (Auburn University), Hasan Abbasi (Oak Ridge National Laboratories), Bin Wang (Auburn University), Norbert Podhorszki (Oak Ridge National Laboratories), Ray Grout (National Renewable Energy Laboratory), Matt Wolf (Georgia Institute of Technology)

Accepted Posters (based on the call for posters)

Indexing a Large-Scale Database of Astronomical Objects
Bin Fu, Eugene Fink, Garth Gibson, Jaime Carbonell (Carnegie Mellon University)

Resource Management for Dynamic MapReduce Clusters within a Multicluster System
Bogdan Ghiț, Nezih Yigitbasi, Dick Epema (Delft University of Technology)

ERMS: Elastic Replication Management System for HDFS
Zhendong Cheng, Zhongzhi Luan, Depei Qian (Sino-German Joint Software Institute (JSI), Beihang University), Alain Roy (University of Wisconsin-Madison)

Dynamic Block-level Cache Management for Cloud Computing Systems
Dulcardo Arteaga, Douglas Otstott, Ming Zhao (Florida International University)

Orestes: a REST protocol for horizontally scalable cloud database access
Felix Gessert, Florian Bücklers, Norbert Ritter (University of Hamburg)

FlexIO: Location-flexible Execution of In Situ Data Analytics for Large Scale Scientific Applications
Fang Zheng, Hongbo Zou, Greg Eisenhauer, Karsten Schwan, Matthew Wolf, Jai Dayal, Tuan-Anh Nguyen, Jianting Cao (Georgia Institute of Technology), Hasan Abbasi, Scott Klasky, Norbert Podhorszki (Oak Ridge National Laboratory), Hongfeng Yu (Sandia National Laboratory, Livermore)

D2T: Doubly Distributed Transactions for High Performance and Distributed Computing
Jai Dayal (Georgia Institute of Technology), Jay Lofstead (Sandia National Labs), Karsten Schwan (Georgia Institute of Technology), Ron Oldfield (Sandia National Labs)

A Performance Comparison of OpenCL and OpenMP for Multi-core CPUs
Jie Shen, Jianbin Fang, Ana Lucia Varbanescu, Henk Sips (Delft University of Technology)

Optimizing Security for Virtual Machine Applications
Naod Duga Jebessa, Guido van 't Noordende, Cees de Laat (University of Amsterdam)

On-Demand Buffers For Large Data Streams In High Performance Computing Clusters
Philip Pum (University of Applied Sciences Upper Austria)

V-cloud: A Peer-to-peer Video Storage-Compute Cloud
Harisankar Haridas, Sriram Kailasam, Prateek Dhawalia, Prateek Shrivastava, Santosh Kumar, Janakiram Dharanipragada (Indian Institute of Technology)

Themis: A Spot-Market Based Automatic Resource Scaling Framework
Stefania Costache (EDF R&D, INRIA), Nikos Parlavantzas, Christine Morin (INRIA), Samuel Kortas (EDF R&D)

Improving SOA Applications Response Time with Service Overload Detection
Valeria Cardellini, Stefano Iannucci (University of Roma "Tor Vergata")

ConPaaS: an Integrated Runtime Environment for Elastic Cloud Applications
Guillaume Pierre, Corina Stratan, Adriana Szekeres, Ana Oprescu, Kaveh Razavi, Thilo Kielmann (Vrije Universiteit Amsterdam), Thorsten Schütt, Jan Stender (Zuse Intitut Berlin), Matej Artač, Aleš Černivec (XLAB)

Parallelization and Distribution for Large Scale Graph Processing
Alexandru Iosup, Nefeli Papapetrou Lampraki, Ate Penders, Marcin Biczak, Yong Guo, Ana Lucia Varbanescu (Delft University of Technology)

Personality Types and Their Parallel Programming Styles
Vincent Hindriksen (StreamComputing)