IEEE Cluster 2016

The Institute of Electrical and Electronics Engineers Cluster 2016

2016.09.13 – 15

PROGRAM BOOK

Taipei, Taiwan

TABLE OF CONTENTS

Message from the General Co-Chairs

On behalf of the entire organizing committee, we wish to welcome you to the IEEE's CLUSTER 2016 conference to be held in Taipei, Taiwan, from September 13 to 15. Since its inauguration in 1999 at Australia, IEEE CLUSTER has since rotated between US, Europe, and Asia. This year's CLUSTER is its fourth time in Asia and the first time in Taiwan.

Over the years, CLUSTER has established itself as a premier conference on cluster computing and served as a major international forum for presenting and sharing recent accomplishments and technological developments in the field of cluster computing as well as the use of cluster systems for scientific and commercial applications. Following this tradition, the organizing committee of CLUSTER 2016 has worked diligently and prepared an excellent and rich program, which includes high quality research papers in 14 technical sessions, together with keynote speeches, panels, industry sessions, workshops, posters, and student mentoring program.

We would like to thank the Program Chair, Bronis R. de Supinski, together with the diligent Program Track Chairs, Kengo Nakjima, Felix Wolf, David Abramson, and Rob Ross, as well as all the program committee members, for putting together a strong technical program. We are also grateful to Torsten Hoefler for organizing the workshops and tutorials, Ewa Deelman and Maciej Malawski for preparing an exciting poster program, Luc Bougé for his dedication and commitment to the student mentoring program, and Mary Hall for organizing the very inspiring panel.

We are indebted to Michela Becchi for her active role in organizing the conference, from proceedings publication to student traveling. We would also like to express our gratitude to the Publicity Co-Chairs, Pavan Balaji, Felix Wolf, Naoya Maruyama, James Lin, and Pi-Cheng Hsiu for reaching out to the community at different parts of the world. We also acknowledge the hard works and dedications by the Local Arrangement Chair, Robert Hsu, the Web Chair, Jerry Chou, and the Finance Chair, Che-Rung Lee.

Our gratitude also goes to the sponsors for their generous supports in organizing CLUSTER 2016 and the many authors who submitted papers. The conference will not be successful without their involvement and participation. It is our great honor to have your attendance, wish you have an enjoyable stay in Taipei.

Jack Dongarra, University of Tennessee, USA
Satoshi Matsuoka, Tokyo Technical University, Japan
Chung-Ta King, National Tsing Hua University, Taiwan

Message from the Program Chair

Cluster computing has become the most widely adopted form of parallel computing since its inception in the early 1990s. It includes the largest systems in the world, such as China's 125+ Petaflops TiahuLight. Cluster systems continue to grow larger and to be applied to a range of new applications, leading to continuing challenges. The research topics include the continuing hardware evolution that enables larger and more powerful systems, the system software and programming systems that can exploit these hardware innovations, the storage systems that can accommodate the increasing production of results enabled by these innovations and, most importantly, the design and continued scaling of applications that exploit the systems that result from the responses to the first three challenges.

The IEEE Cluster conference series is a major venue in which to present the successful efforts to address the challenges presented by cluster computing. This year continues several changes made for IEEE Cluster 2015. In particular, we have four submission areas that directly reflect the major challenges: Applications, Algorithms, and Libraries (chair: Kengo Nakajima, the University of Tokyo); Architecture, Networks/Communication, and Management (chair: Felix Wolf, Technische Universität Darmstadt); Programming and System Software (chair: David Abramson, University of Queensland); and Data, Storage, and Visualization (chair: Rob Ross, Argonne National Laboratory), with the authors being asked to provide their first and second area preferences. The program committee consists of top-quality researchers in these areas with a balance in geographic and other demographics, which ensured that the reviews reflect the latest achievements in cluster computing.

Overall, Cluster 2016 received a record 162 paper submissions, up slightly from 2015 and continuing the growth compared to prior years. Minor reshuffling provided nearly perfect balance across the areas, with approximately 40 papers in each. All papers were rigorously reviewed, with all papers receiving at least three and many four or more reviews with substantive comments. After an on-line discussion process, we accepted 39 full papers (24.1% acceptance rate) and 16 short papers, with a short paper being given a 4-page publication space in the IEEE Digital Library and a 15-minute presentation slot during the conference, plus a note that the full & extended version of the paper would be acceptable for other conferences or journals, just as the evolved versions of a conference paper later being accepted in a journal. From the full papers, we have identified four papers that are Best Paper Award Nominees, which will be presented in a special session and from which an overall Best Paper will be selected by a panel of judges during the conference. Combined with other program elements including the keynotes, posters and workshops, as well as the student mentoring program, IEEE Cluster 2016 continues the level of technical excellence that recent IEEE Cluster conferences have established and upon which we expect future ones to expand.

This level of technical achievement would not be possible without the invaluable efforts of many others. My sincere appreciation is extended first to the area chairs, who made my role easy. I also thank the many paper program committee members, as well as their subreviewers, who have contributed many hours in their reviews and discussions without which we could not realize our vision of technical excellence. Further, I thank the Cluster 2015 conference committee and Mark Montague of Linklings, who have provided invaluable assistance in the paper review process and various other places that a successful conference requires. Finally, and most of all, the entire committee acknowledges the contributions of the authors who submitted their high quality work, for without community support the conference would not happen. We look forward to the excellent program in Taipei in September and to seeing you there.

Bronis R. de Supinski, Lawrence Livermore national Laboratory, USA

Acknowledgement

Organizers

National Tsing Hua University

National Center for High-performance Computing

Research Center for Information Technology Innovation

Co-Organizer

Taiwan Association of Cloud Computing

Sponsors

IEEE

IEEE Computer Society

MediaTek

National Science Foundation

Nvidia

TeraSoft Inc.

Government Supports

Bureau of Foreign Trade

Department of Information and Tourism, Taipei City Government

Ministry of Education, R.O.C. (Taiwan)

Ministry of Foreign Affairs, R.O.C. (Taiwan)

Ministry of Science and Technology

Organizing Committees

General Chair	Jack Dongarra	University of Tennessee, USA
	Satoshi Matsuoka	Tokyo Technical University, Japan
	Chung-Ta King	National Tsing Hua University, Taiwan

Technical Program

Program Chair	Bronis R. de Supinski	Lawrence Livermore National Laboratory, USA
Workshop/Tutorial Chair	Torsten Hoefler	ETH Zürich University, Switzerland
Panels Chair	Mary Hall	University of Utah, USA
Poster Chair	Ewa Deelman	University of Southern California, USA
Poster Vice Chair	Maciej Malawski	AGH University of Science and Technology, Poland
Proceedings Chair	Michela Becchi	University of Missouri – Columbia, USA
Student Mentoring Chair	Luc Bougé	ENS Rennes, IRISA/INRIA, France
Student Travel Chair	Michela Taufer	University of Delaware, USA

Logistics

Publicity Co-chairs	Pavan Balaji	Argonne National Laboratory, USA
	Pi-Cheng Hsiu	Research Center for Information Technology Innovation, Taiwan
	James Lin	Shanghai Jiao Tong University, China
	Naoya Maruyama	RIKEN, Japan
	Felix Wolf	Technische Universität Darmstadt, Germany
Local Arrangement Chair	Robert Hsu	Chung Hua University, Taiwan
Finance Chair	Che-Rung Lee	National Tsing Hua University, Taiwan

Technical Program Committee Chairs

Applications, Algorithms, and Libraries		
Chair	Kengo Nakjima	The University of Tokyo, Japan
Architecture, Networks/Communication, and Managements		
Chair	Felix Wolf	Technische Universität Darmstadt, Germany
Programming and System Software		
Chair	David Abramson	University of Queensland, Australia
Data, Storage, and Visualization		
Chair	Rob Ross	Argonne National Laboratory, USA

Technical Program Committees

Applications, Algorithms, and Libraries	
Wolfgang Bangerth	Texas A&M, USA
Costas Bekas	IBM Research, Zurich, Switzerland
Carsten Burstedde	University of Bonn, Germany
Xing Cai	Simula Laboratory, Norway
Christophe Calvin	CEA, France
Edmond Chow	Georgia Tech., USA
Karen Devine	Sandia National Laboratories, USA
Tony Drummond	Lawrence Berkeley National Laboratory, USA
David Eder	Lawrence Livermore National Laboratory, USA
Laura Grigori	INRIA, France

Takeshi Iwashita	Hokkaido University, Japan
Zhong Jin	Chinese Academy of Science, China
Jakub Kurzak	University of Tennessee, USA
Osni Marques	Lawrence Berkeley National Laboratory, USA
Amanda Peters-Randles	Duke University, USA
Siva Rajamanickam	Sandia National Laboratories, USA
Sabela Ramos	Universidade da Coruna, Spain
Hari Sundar	University of Utah, USA
Weichung Wang	National Taiwan University, Taiwan
Gerhard Wellein	University of Erlangen-Nuremberg, Germany
Rio Yokota	Tokyo Institute of Technology, Japan

Architecture, Networks/Communication, and Managements

Abhinav Bhatele	Lawrence Livermore National Laboratory, USA
Ron Brightwell	Sandia National Laboratories, USA
Michael Brim	Oak Ridge National Laboratory, USA
Evelyn Duesterwald	IBM Corporation, USA
Holger Fröning	University of Heidelberg, Germany
Todd Gamblin	Lawrence Livermore National Laboratory, USA
Markus Geimer	Jülich Supercomputing Centre, Germany
Brice Goglin	INRIA, France
Wei Huang	AMD, USA
Sascha Hunold	TU Vienna, Austria
Stefanos Kaxiras	Uppsala University, Sweden
Michihiro Koibuchi	National Institute of Informatics, Japan
Masaaki Kondo	University of Tokyo, Japan
James Lin	Shanghai Jiao Tong University, China
Naoya Maruyama	RIKEN, Japan
Dimitrios Nikolopoulos	Queen's University of Belfast, United Kingdom
Miquel Pericas	Chalmers University of Technology, Sweden
Phil Roth	Oak Ridge National Laboratory, USA
Shinji Sumimoto	Fujitsu Labs, Japan
Keith Underwood	Intel Corporation, USA
Osman Unsal	Barcelona Supercomputing Center, Spain
Ana Varbanescu	University of Amsterdam, Netherland

Data, Storage, and Visualization

Gabriel Antoniu	INRIA Rennes, France
Andre Brinkmann	Johannes Gutenberg-Universitaet Mainz, Germany
Yong Chen	Texas Tech University, USA
Matthew Curry	Sandia National Laboratories, USA
Qiang Guan	Los Alamos National Laboratory, USA
Shadi Ibrahim	INRIA Rennes, France
Florin Isaila	Universidad Carlos III de Madrid, Spain
Hai Jin	Huazhong University of Science and Technology, China
Julian Kunkel	DKRZ, Germany
Qing Liu	Oak Ridge National Laboratory, USA
Xiaosong Ma	Qatar Foundation, Qatar
Carlos Maltzahn	University of California, Santa Cruz, USA
Ken Moreland	Sandia National Laboratories, USA
Jorji Nonaka	RIKEN, Japan

Manish Parashar	Rutgers University, USA
Maria S. Perez	Universidad Politecnica de Madrid, Spain
Thomas Peterka	Argonne National Laboratory, USA
Brad Settlemyer	Los Alamos National Laboratory, USA
Haiying Shen	Clemson University, USA
Han-Wei Shen	The Ohio State University, USA
Osamu Tatebe	University of Tsukuba, Japan
Nick Wright	Lawrence Berkeley National Laboratory, USA

Programming and System Software	
Franck Cappello	Argonne National Laboratory, USA
Ewa Deelman	ISI, USC, USA
Luiz DeRose	Cray, Inc. , USA
William Gropp	University of Illinois Urbana-Champaign, USA
Elizabeth Jessup	University of Colorado, USA
Thilo Kielmann	Vrije Universiteit, Netherland
Dieter Kranzlmueller	Ludwig-Maximilians-Universitaet, Munich, Germany
Zhiling Lan	Illinois Institute of Technology, USA
Dong Li	University of California, Merced, USA
Hatem Ltaief	KAUST Supercomputing Laboratory, Saudi Arabia
Abid Malik	University of Houston, USA
Satoshi Matsuoka	Tokyo Institute of Technology, Japan
Raymond Namyst	Universite de Bordeaux, France
Dhabaleswar Panda	Ohio State University, USA
Jelena Pjesivac-Grbovic	Google, USA
Ioan Raicu	Illinois Institute of Technology, USA
Kento Sato	Lawrence Livermore National Laboratory, USA
John Shalf	Lawrence Berkeley National Laboratory, USA
Michela Taufer	University of Delaware, USA
Kenjiro Taura	University of Tokyo, Japan
Greg Watson	Oak Ridge National Labs, USA
Steven Young	University of Oxford, United Kingdom

Posters Committees

Rosa Filgueira	Edinburgh University, United Kingdom
Tristan Glatard	McGill University, Canada
Soonwook Hwang	Korea Institute of Science and Technology, Korea
Ming Jiang	Lawrence Livermore National Laboratory, USA
Dieter Kranzlmueller	Ludwig Maximilian University of Munich, Germany
Dariusz Krol	AGH University of Science and Technology, Poland
Che-Rung Lee	National Tsing Hua University, Taiwan
Jaroslaw Nabrzyski	University of Notre Dame, USA
Maria S. Perez	Technical University of Madrid, Spain
Tom Peterka	Argonne National Laboratory, USA
Seetharami Seelam	IBM, USA
Satoshi Sekiguchi	National Institute of Advanced Industrial Science and Technology, Japan
Frederic Sutter	École Normale Supérieure de Lyon, France
Michela Taufer	University of Delaware, USA

Program at a Glance

DATE	Day 1 - Sep. 13 (Tue.)			Day 2 - Sep. 14 (Wed.)			Day 3 - Sep. 15 (Thu.)		
Venue	Grand Hall	Grand Hall A	Grand Hall B	Grand Hall	Grand Hall A	Grand Hall B	Grand Hall	Grand Hall A	Grand Hall B
09:00	Opening Session			Introduce Cluster 2017		Conference Posters	Keynote 3		Conference Posters
09:10–09:50	Keynote 1			Keynote 2					
10:00–10:20	Coffee Break			Coffee Break			Coffee Break		
10:30–10:50	Best Paper Nominees (Paper Session 1)			Combined Paper/Panel Session (Paper Session 5)			Paper Session 9	Paper Session 10	Poster Removal
11:00–11:50									
12:00–12:20									Lunch
12:30–12:50	Lunch	Student Mentoring Program		Lunch	Student Mentoring Program		Lunch	Student Mentoring Program	
13:00–13:50									
14:00–14:50	Paper Session 2	Paper Session 3		Paper Session 6	Paper Session 7		Paper Session 11	Paper Session 12	Workshop FTS
15:00–15:20			Poster Setup						
15:30–15:50	Coffee Break			Coffee Break			Coffee Break		
16:00–16:50	Paper Session 4	Tutorials State-of-the Art in Slurm, MPI-PGAS and BigData		Paper Session 8	Industry Session		Paper Session 13	Paper Session 14	Workshop FTS
17:00–17:50				Panel Session					
18:00–18:20									
18:30–19:50	Welcome Reception & Poster Reception (Grand Hall+Grand Hall B)			Gala Dinner (Grand Hall)					
20:00–20:50									

Keynote Speakers

Keynote 1 – Theory and Practice in HPC: Modeling, Programming, and Networking

Torsten Hoeffler *(ETH Zürich University)*

Date: September 13, 2016
Time: 09:00-10:10
Room: Grand Hall
Chair: Jack Dongarra *(University of Tennessee, Knoxville)*

Abstract

We advocate the usage of mathematical models and abstractions in practical high-performance computing. For this, we show a series of examples and use-cases where the abstractions introduced by performance models can lead to clearer pictures of the core problems and often provide non-obvious insights. We start with models of parallel algorithms leading to close-to-optimal practical implementations. We continue our tour with distributed-memory programming models that provide various abstractions to application developers. A short digression on how to measure parallel systems shows common pitfalls of practical performance modeling. Application performance models based on such accurate measurements support insight into the resource consumption and scalability of parallel programs on particular architectures. We close with a demonstration of how mathematical models can be used to derive practical network topologies and routing algorithms. In each of these areas, we demonstrate newest developments but also point to open problems. All these examples testify to the value of modeling in practical high-performance computing. We assume that a broader use of these techniques and the development of a solid theory for parallel performance will lead to deep insights at many fronts.

Biography

Torsten Hoeffler is an Assistant Professor of Computer Science at ETH Zürich, Switzerland. Before joining ETH, he led the performance modeling and simulation efforts of parallel petascale applications for the NSF-funded Blue Waters project at NCSA/UIUC. He is also a key member of the Message Passing Interface (MPI) Forum where he chairs the "Collective Operations and Topologies" working group. Torsten won best paper awards at the ACM/IEEE Supercomputing Conference SC10, SC13, SC14, EuroMPI'13, HPDC'15, HPDC'16, IPDPS'15, and other conferences. He published numerous peer-reviewed scientific conference and journal articles and authored chapters of the MPI-2.2 and MPI-3.0 standards. He received the Latsis prize of ETH Zurich as well as an ERC starting grant in 2015. His research interests revolve around the central topic of "Performance-centric System Design" and include scalable networks, parallel programming techniques, and performance modeling. Additional information about Torsten can be found on his homepage at http://htor.inf.ethz.ch.

Keynote 2 – Caches All the Way Down: Infrastructure for Data Science

David Abramson ***(University of Queensland)***

Date: September 14, 2016
Time: 09:00-10:10
Room: Grand Hall
Chair: Satoshi Matsuoka *(Tokyo Institute of Technology)*

Abstract

The rise of big data science has created new demands for modern computer systems. While floating point performance has driven computer architecture and system design for the past few decades, there is renewed interest in the speed at which data can be ingested and processed. Early exemplars such as Gordon, the NSF funded system at the San Diego Supercomputing Centre, shifted the focus from pure floating point performance to memory and IO rates. At the University of Queensland we have continued this trend with the design of FlashLite, a parallel cluster equiped with large amounts of main memory, Flash disk, and a distributed shared memory system (ScaleMP's vSMP). This allows applications to place data "close" to the processor, enhancing processing speeds. Further, we have built a geographically distributed multi-tier hierarchical data fabric called MeDiCI, which provides an abstraction of very large data stores across the metropolitan area. MeDiCI leverages industry solutions such as IBM's Spectrum Scale and SGI's DMF platforms.

Caching underpins both FlashLite and MeDiCI. In this talk I will describe the design decisions and illustrate some early application studies that benefit from the approach.

Biography

Director, Research Computing Centre David has been involved in computer architecture and high performance computing research since 1979. He has held appointments at Griffith University, CSIRO, RMIT and Monash University. Prior to joining UQ, he was the Director of the Monash e-Education Centre, Science Director of the Monash e-Research Centre, and a Professor of Computer Science in the Faculty of Information Technology at Monash. From 2007 to 2011 he was an Australian Research Council Professorial Fellow. David has expertise in High Performance Computing, distributed and parallel computing, computer architecture and software engineering. He has produced in excess of 200 research publications, and some of his work has also been integrated in commercial products. One of these, Nimrod, has been used widely in research and academia globally, and is also available as a commercial product, called EnFuzion, from Axceleon. His world-leading work in parallel debugging is sold and marketed by Cray Inc, one of the world's leading supercomputing vendors, as a product called ccdb. David is a Fellow of the Association for Computing Machinery (ACM), the Institute of Electrical and Electronic Engineers (IEEE), the Australian Academy of Technology and Engineering (ATSE), and the Australian Computer Society (ACS). He is currently a visiting Professor in the Oxford e-Research Centre at the University of Oxford.

Keynote 3 – Who is afraid of I/O? Exploring I/O Challenges and Opportunities at the Exascale

Michela Taufer ***(University of Delaware)***

Date: September 15, 2016
Time: 09:00-10:00
Room: Grand Hall
Chair: Chung-Ta King *(National Tsinghua University)*

Abstract

Clear trends in the past and current petascale systems (i.e., Jaguar and Titan) and the new generation of systems that will transition us toward exascale (i.e., Aurora and Summit) outline how concurrency and peak performance are growing dramatically, however, I/O bandwidth remains stagnant. Next-generation systems are expected to deliver 7 to 10 times higher peak floating-point performance with only 1 to 2 times higher PFS bandwidth compared to the current generation.

Data intensive applications, especially those exhibiting bursty I/O, must take this aspect into consideration and be more selective about what data is written to disk and how the data is written. In addressing the needs of these applications, can we take advantage of a rapidly changing technology landscape, including containerized environments, burst buffers, and in-situ/in-transit analytics? Are these technologies ready to transition these applications to exascale? In general, existing software components managing these technologies are I/O-ignorant, resulting in systems running the data intensive applications that exhibit contentions, hot spots, and poor performance.

In this talk, we explore challenges when dealing with I/O-ignorant high performance computing systems and opportunities for integrating I/O awareness in these systems. Specifically, we present solutions that use I/O awareness to reduce contentions in scheduling policies managing under provisioned systems with burst buffers, and to mitigate data movements in data-intensive simulations. Our proposed solutions go beyond high performance computing and develop opportunities for interdisciplinary collaborations.

Biography

Michela Taufer is an associate professor in the Computer and Information Sciences Department at the University of Delaware. She earned her master's degrees in Computer Engineering from the University of Padova (Italy) and her doctoral degree in Computer Science from the Swiss Federal Institute of Technology (Switzerland). From 2003 to 2004 she was a La Jolla Interfaces in Science Training Program (LJIS) Postdoctoral Fellow at the University of California San Diego (UCSD) and The Scripps Research Institute (TSRI), where she worked on interdisciplinary projects in computer systems and computational chemistry. From 2005 to 2007, she was an Assistant Professor at the Computer Science Department of the University of Texas at El Paso (UTEP). She joined the University of Delaware in 2007 as an Assistant Professor and was promoted to Associate Professor with tenure in 2012. Taufer's research interests include scientific applications and their advanced programmability in heterogeneous computing (i.e., multi-core and many-core platforms, GPUs); performance analysis, modeling, and optimization of multi-scale applications on heterogeneous computing, cloud computing, and volunteer computing; numerical reproducibility and stability of large-scale simulations on multi-core platforms; big data analytics and MapReduce.

Industry Session Speaker

NVIDIA Industry Session – HPC for AI Development

Date: September 14, 2016
Time: 16:00-17:00
Room: Grand Hall A
Chair: Jerry Chou *(National Tsinghua University)*

Abstract

Artificial Intelligence has recently taken on a new life. It is being applied or used in robotics, drone, chatbot, facial recognition and many other applications. Many large enterprises and institution have invested and continue to invest in the development of artificial intelligence technology. Many new algorithms are being developed and many of them are compute intensive. In this talk, we will discuss how HPC is being use to support the development of artificial intelligence development.

Biography

Director and Chief Solution Architect NVIDIA Inc. Asia Pacific Chief Scientific Computing Officer and Professor Shanghai Jiao Tong University Visiting Professor and Chief Scientific Computing Advisor BGI Prof Simon See is currently the High Performance Computing Technology Director and Chief Solution Architect for NVIDIA Inc, Asia and also a Professor and Chief Scientific Computing Officer in Shanghai Jiao Tong University. Concurrently Prof See is also the Chief Scientific Computing Advisor for BGI (China). Currently Prof See has an position in NanyangTechnological Univ (Singapore) and King-Mong Kung Univ of Technology (Thailand) . His research interests are in the area of high performance computing, artificial intelligence, machine learning , computational science, applied mathematics and simulation methodology. He has published over 100 papers in these areas and has won various awards. Prof See is also member of SIAM, IEEE and IEE. He has also been a committee member of more than 50 conferences. Dr. See graduated from University of Salford (UK) with a Ph.D. in electrical engineering and numerical analysis in 1993. Prior to joining NVIDIA, Dr See worked for SGI, DSO National Lab. of Singapore, IBM, International Simulation Ltd (UK), Sun Microsystems and Oracle. He is also providing consultancy to a number of national research and supercomputing centers.

Daily Program

September 13, 2016 (Tue.)

Opening Remarks

Time: 09:00-09:10
Room: Grand Hall

Keynote 1

Time: 09:10-10:10
Room: Grand Hall
Chair: Jack Dongarra *(University of Tennessee, Knoxville)*

09:10-10:10 **Theory and Practice in HPC: Modeling, Programming, and Networking**
Torsten Hoeffler *(ETH Zürich University)*

Coffee Break
Time: 10:10-10:30

Paper Session 1 - Best Paper Nominees

Time: 10:30-12:30
Room: Grand Hall
Chair: Bronis R. de Supinski *(Lawrence Livermore National Laboratory)*

10:30-11:00 **Adaptive and Dynamic Design for MPI Tag Matching**
Mohammadreza Bayatpour, Hari Subramoni, Sourav Chakraborty, and Dhabaleswar K. Panda *(The Ohio State University)*

11:00-11:30 **A Low Disk-Bound Transaction Logging System for In-Memory Distributed Data Stores**
Dayal Dilli and Kenneth B. Kent *(University of New Brunswick)* and Yang Wang and Chengzhong Xu *(Chinese Academy of Science)*

11:30-12:00 **Realizing Out-of-Core Stencil Computations using Multi-Tier Memory Hierarchy on GPGPU Clusters**
Toshio Endo *(Tokyo Institute of Technology)*

12:00-12:30 **Parallel DTFE Surface Density Field Reconstruction**
Esteban Rangel *(Northwestern University)*; Nan Li, Salman Habib, and Tom Peterka *(Argonne National Laboratory)*; and Ankit Agrawal, Wei-keng Liao, and Alok Chowdhary *(Northwestern University)*

Lunch Break
Time: 12:30-14:00

Student Mentoring Program

Time: 12:30-14:00
Room: Grand Hall A

Paper Session 2- Network Performance

Time: 14:00-15:30
Room: Grand Hall
Chair: Pavan Balaji *(Argonne National Laboratory)*

14:00-14:30 Evaluation of Topology-Aware Broadcast Algorithms for Dragonfly Networks
Matthieu Dorier, Misbah Mubarak, and Rob Ross *(Argonne National Laboratory)*; Jianping Kelvin Li *(University of California, Davis)*; Christopher D. Carothers *(Computer Science Department, Rensselaer Polytechnic Institute)*; and Kwan-Liu Ma *(University of California, Davis)*

14:30-15:00 (SAI) Stalled, Active and Idle: Characterizing Power and Performance of Large-Scale Dragonfly Networks
Taylor Groves, Ryan Grant, Scott Hemmert, Simon Hammond, and Michael Legenhagen *(Sandia National Labs)* and Dorian Arnold *(University of New Mexico)*

15:00-15:30 Compiler-Assisted Overlapping of Communication and Computation in MPI Applications
Jichi Guo and Qing Yi *(University of Colorado - Colorado Springs)* and Jiayuan Meng, Junchao Zhang, and Pavan Balaji *(Argonne National Laboratory)*

Paper Session 3 - Virtualization

Time: 14:00-15:30
Room: Grand Hall A
Chair: Osamu Tatebe *(University of Tsukuba)*

14:00-14:30 vProbe: Scheduling Virtual Machines on NUMA Systems
Song Wu, Huahua Sun, Like Zhou, Qingtian Gan, and Hai Jin *(Huazhong University of Science and Technology)*

14:30-15:00 GLAP: Distributed Dynamic Workload Consolidation through Gossip-Based Learning
Mansour Khelghatdoust and Vincent Gramoli *(Data61-CSIRO / The University of Sydney)* and Daniel Sun *(Data61-CSIRO)*

15:00-15:30 CORP: Cooperative Opportunistic Resource Provisioning for Short-Lived Jobs in Cloud Systems
Jinwei Liu, Haiying Shen, and Liuhua Chen *(Clemson University)*

Coffee Break
Time: 15:30-16:00

Paper Session 4 - Memory Optimization

Time: 16:00-17:30
Room: Grand Hall
Chair: Mary Hall *(University of Utah)*

16:00-16:30 Towards Resource Disaggregation - Memory Scavenging for Scientific Workloads
Alexandru Uta, Ana-Maria Oprescu, and Thilo Kielmann *(VU University Amsterdam)*

16:30-17:00 CHOPPER: Optimizing Data Partitioning for In-Memory Data Analytics Frameworks
Arnab Kumar Paul, Wenjie Zhuang, and Luna Xu *(Virginia Tech);* Min Li *(IBM Almaden Research);* Mustafa Rafique *(IBM Research-Ireland);* and Ali R. Butt *(Virginia Tech)*

17:00-17:30 Improving Collective I/O Performance Using Non-Volatile Memory Devices
Giuseppe Congiu and Sai Narasimhamurthy *(seagate)* and TIm Seuss and André Brinkmann *(University of Mainz)*

Tutorials - State-of-the Art in Slurm, MPI-PGAS and BigData

Time: 16:00-18:00
Room: Grand Hall A

16:00-16:40 **Resource and Job Management on HPC Clusters with Slurm: Administration, Usage and Performance Evaluation**
Yiannis Georgiou *(BULL)*

16:40-17:20 **PGAS and Hybrid MPI+PGAS Programming Models on Modern HPC Clusters with Accelerators**
Dhabaleswar K. (DK) Panda *(The Ohio State University)*

17:20-18:00 **Big Data Meets HPC: Exploiting HPC Technologies for Accelerating Apache Hadoop, Spark and Memcached**
Dhabaleswar K. (DK) Panda *(The Ohio State University)*

Welcome Reception & Poster Exhibition
Time: 18:30-20:00
Room: Grand Hall+Grand Hall B
Chair: Ewa Deelman *(University of Southern California Information Sciences Institute)*

PO-1 **SMARTPARTITION: Efficient Partitioning for Natural Graphs**
Chengfei Zhang, Yiming Zhang, Dongsheng Li, Jia Li, and Minne Li *(National University of Defense Technology)*

PO-2 **MBL: A Multi-Stage Bufferless High-radix Router**
Wenxiang Yang *(National University of Defense Technology)*

PO-3 **Acclerating I/O performance of SVM on HDFS**
Mao Ye and Jun Wang *(University of Central Florida)*

PO-4 **Energy and Performance Efficient Underloading Detection Algorithm of Virtual Machines in Cloud Data Centers**
Juan Fang, Lifu Zhou, Xiaoting Hao, Min Cai, and Xingtian Ren *(Beijing University of Technology)*

PO-5 **Machine Status Prediction for Dynamic and Heterogenous Cloud Environment**
Jinliang Xu, Ao Zhou, Shangguang Wang, Qibo Sun, Jinglin Li, and Fangchun Yang *(Beijing University of Posts and Telecommunications)*

PO-6 **Design and Analysis of Fault Tolerance Mechanisms for Big Data Transfers**
Preethika Kasu *(Ajou University)*, Youngjae Kim and Sungyong Park *(Sogang University)*, and Scott Atchley and Geoffroy R. Vallee *(Oak Ridge National Laboratory)*

PO-7 **Time Optimization Modeling for Big Data Placement and Analysis for Geo-Distributed Data Centers**
Awais Khan, Muhammad Attique, and Tae-Sun Chung *(Ajou University)* and Youngjae Kim *(Sogang University)*

PO-8 **Enhancing Performance of Large-scale Electronic Structure Calculations with Many-core Computing**
Hoon Ryu and Yosang Jeong *(Korea Institute of Science and Technology Information)*

PO-9 **Reduced-Precision Floating-Point Formats on GPUs for High Performance and Energy Efficient Computation**
Daichi Mukunoki and Toshiyuki Imamura *(RIKEN AICS)*

PO-10 **Themis: A Scalable Performance Evaluation Framework for Virtualized Datacenter**
Zhengmin Li *(CNCERT/CC)*; Di Zhang *(ICT, CAS)*; Xinran Liu *(CNCERT/CC*); Bin Sun *(Beijing University of Posts and Telecommunications)*; and Zhicheng Yao and Xiufeng Sui *(ICT, CAS)*

PO-11 **Conflict Prediction-based Transaction Execution for Transactional Memory in Multi-Core In-Memory Databases**
Min Yoon, Moon-Hwan Kang, Yeon-Woo Jang, and Jae-Woo Chang *(Chonbuk National University)*

PO-12 **Skyline Service Selection based on QoS Prediction**
Yan Guo and Shangguang Wang *(Beijing University of Posts and Telecommunications)*

PO-13 **Minimizing CMT Miss Penalty in Selective Page-level Address Mapping Table**
Ronnie Mativenga and Joon-Young Paik *(Ajou University)*; Junghee Lee *(University of Texas, San Antonio)*; Tae-Sun Chung *(Ajou University)*; and Youngjae Kim *(Sogang University)*

PO-14 **Efficient Semantic-Aware Coflow Scheduling for Data-Parallel Jobs**
Ziyang Li, Yiming Zhang, Yunxiang Zhao, and Dongsheng Li *(National University of Defense Technology)*

PO-15 **Quick Eviction of Virtual Machines Through Proactive Snapshots**
Dinuni Fernando, Hardik Bagdi, Yaohui Hu, Ping Yang, and Kartik Gopalan *(Binghamton University)* and Charles Kamhoua and Kevin Kwiat *(Air Force Rome Research Laboratory)*

PO-16 **Spatial Locality Aware, Fast, and Scalable SLINK Algorithm for Commodity Clusters**
Poonam Goyal, Sonal Kumari, Sumit Sharma, Vivek Choudhary, Navneet Goyal, and Sundar Balasubramaniam (Birla Institute of Technology and Science)

PO-17 **High Throughput Log-based Replication for Many Small In-memory Objects**
Kevin Beineke *(Heinrich-Heine Universitaet Duesseldorf - Institut fuer Informatik, Abteilung Betriebssysteme)*

PO-18 **TwinPCG: Dual Thread Redundancy with Forward Recovery for Conjugate Gradient Methods**
Kiril Dichev and Dimitrios Nikolopoulos *(Queen's University Belfast)*

PO-19 **Optimizing Locality by Topology-aware Placement for a Task Based Programming Model**
Jens Gustedt, Emmanuel Jeannot, and Farouk Mansouri *(INRIA)*

PO-20 **SSDUP: An Efficient SSD Write Buffer Using Pipeline**
Ming Li, Xuanhua Shi, Wei Liu, and Hai Jin *(Huazhong University of Science and Technology)* and Yong Chen *(Texas Tech University)*

PO-21 **StreamingRPHash: Random Projection Clustering of High-Dimensional Data in a MapReduce Framework**
Jacob Franklin, Samuel Wenke, Sadiq Quasem, Lee Carraher, and Philip Wilsey *(University of Cincinnati)*

PO-22 **SuperGlue: Standardizing Glue Components for HPC Workflows**
Jay Lofstead (Sandia National Laboratories) and Alexis Champsaur, Jai Dayal, Matthew Wolf, and Greg Eisenhauer (Georgia Institute of Technology)

September 14, 2016 (Wed.)

Presentation of Cluster 2017

Time: 09:00-09:10
Room: Grand Hall

Keynote 2

Time: 09:10-10:10
Room: Grand Hall
Chair: Satoshi Matsuoka *(Tokyo Institute of Technology)*

09:10-10:10 **Caches All the Way Down: Infrastructure for Data Science**
David Abramson *(University of Queensland)*

Coffee Break
Time: 10:10-10:30

Paper Session 5/Panel Session - Large-Scale System Analysis and Modeling: Combined Paper

Time: 10:30-12:30
Room: Grand Hall
Chair: Todd Gamblin *(Lawrence Livermore National Laboratory)*

This talk will provide a broad overview of analysis and modeling for large-scale systems, starting with the behavior of single applications and benchmarks, and proceeding to the analysis of entire HPC systems or data centers. Increasingly, the performance of applications on parallel systems depends not only on an application's own performance, but also on its inputs and the performance, resource usage, and scaling behavior of other applications. Worse, performance variability resulting from resource contention makes it hard to model the performance of a single application without detailed measurements from an entire cluster or data center. This talk will discuss challenges facing performance analysts on modern systems, as well as some early solutions for using machine learning to understand application performance in a variable computing environment.

10:30-11:00 **Topic Overview: Large-Scale System Analysis and Modeling**
Todd Gamblin *(Lawrence Livermore National Laboratory)*
This talk will provide a broad overview of the topic of large-scale analysis and modeling. The talk will provide the context for the four papers included in this session and how those papers fit into that context.

11:00-11:12 **Fast Multi-Parameter Performance Modeling**
Alexandru Calotoiu *(TU Darmstadt)*, David Beckingsale and Christopher William Earl *(Lawrence Livermore National Laboratory)*, Torsten Hoefler *(ETH Zurich)*, Ian Karlin and Martin Schulz *(Lawrence Livermore National Laboratory)*, and Felix Wolf *(TU Darmstadt)*

11:12-11:24 **Active Learning in Performance Analysis**
Dmitry Duplyakin and Jed Brown *(University of Colorado)* and Robert Ricci *(University of Utah)*

11:24-11:36 **A Model for Weak Scaling to Many Gpus at The Basis of The Linpack Benchmark**
David Rohr and Jan de Cuveland *(Frankfurt Institute for Advanced Studies)* and Volker Lindenstruth *(Frankfurt Institute for Advanced Studies, Goethe Universität Frankfurt)*

11:36-11:48 **When Amdahl Meets Young/Daly**
Aurélien Cavelan *(ENS Lyon, Inria)*; Jiafan Li *(East China Normal University)*; Yves Robert *(ENS Lyon, Inria)*; and Hongyang Sun *(ENS de Lyon, Inria)*

11:48-12:30 **Author Panel: Large-Scale System Analysis and Modeling**

Moderator: Todd Gamblin *(Lawrence Livermore National Laboratory)*

Panelists: Felix Wolf *(TU Darmstadt)*; Dmitry Duplyakin *(University of Colorado)*; David Rohr *(Frankfurt Institute for Advanced Studies)*; and Hongyang Sun *(ENS de Lyon, Inria)*
Authors of papers included in this session will answer questions about the general topic of large-scale system analysis and modeling as well as about their papers.

Lunch Break
Time: 12:30-14:00

Student Mentoring Program
Time: 12:30-14:00
Room: Grand Hall A

Paper Session 6 - Energy and Resilience
Time: 14:00-15:30
Room: Grand Hall
Chair: David Abramson *(University of Queensland)*

14:00-14:15 **A Case for Criticality Models in Exascale Systems**
Brian Kocoloski *(University of Pittsburgh)*; Leonardo Piga, Wei Huang, and Indrani Paul *(AMD)*; and John Lange *(University of Pittsburgh)*

14:15-14:30 **Dynamically Building Energy Proportional Data Centers with Heterogeneous Computing Resources**
Violaine Villebonnet *(Inria Avalon LIP ,ENS Lyon, University of Lyon - IRIT, University of Toulouse)*; Georges Da Costa *(IRIT, University of Toulouse)*; Laurent Lefevre *(Inria Avalon LIP ,ENS Lyon, University of Lyon)*; and Jean-Marc Pierson and Patricia Stolf *(IRIT, University of Toulouse)*

14:30-15:00 **On Energy Proportionality and Time-energy Performance of Heterogeneous Clusters**
Lavanya Ramapantulu, Dumitrel Loghin, and Yong Meng Teo *(National University of Singapore)*

15:00-15:30 **Unequal Failure Protection Coding Technology for Cloud Storage Systems**
Yupeng Hu *(Hunan University)*, Yonghe Liu *(UT-Arlington)*, Wenjia Li *(New York Institute of Technology)*, Nong Xiao *(National University of Defense Technology)*, and Zheng Qin and Shu Yin *(Hunan University)*

Paper Session 7 - Resource Management
Time: 14:00-15:30
Room: Grand Hall A
Chair: Ewa Deelman *(University of Southern California Information Sciences Institute)*

14:00-14:30 **Probabilistic Network-Aware Task Placement for MapReduce Scheduling**
Haiyeng Shen and Ankur Sarker *(Clemson University)*, Lei Yu *(Georgia Tech)*, and Feng Deng *(Clemson University)*

14:30-14:45 **Tiresias: Low-overhead Sample Based Scheduling with Task Hopping**
Chunliang Hao *(Institute of Software, Chinese Academy of Science)*; Jie Shen *(Imperial College London)*; and Heng Zhang *(Institute of Software, Chinese Academy of Science)*

14:45-15:00 **Machine Learning Predictions of Runtime and IO Traffic on High-End Clusters**
Ryan McKenna and Stephen N. Herbein *(University of Delaware)*, Adam Moody and Todd Gamblin *(Lawrence Livermore National Laboratory)*, and Michela Taufer *(University of Delaware)*

15:00-15:30 **Exploring Plan-Based Scheduling for Large-Scale Computing Systems**
Xingwu Zheng, Zhou Zhou, Xu Yang, Zhilin Lan, and Jia Wang *(Illinois Institute of Technology)*

Coffee Break
Time: 15:30-16:00

Paper Session 8 - In-Situ Applications (Paper Session 8)
Time: 16:00-17:00
Room: Grand Hall
Chair: Michela Taufer *(University of Delaware)*

16:00-16:30 **Adaptive Performance-Constrained In Situ Visualization of Atmospheric Simulations**
Matthieu Dorier *(Argonne National Laboratory)*; Robert Sisneros *(NCSA, UIUC)*; Leonardo Bautista Gomez and Tom Peterka *(Argonne National Laboratory)*; Leigh Orf *(UW - Madison)*; Lokman Rahmani *(ENS Rennes)*; Gabriel Antoniu *(Inria Rennes Bretagne Atlantique)*; and Luc Bougé *(ENS Rennes)*

16:30-17:00 **Bredala: Semantic Data Redistribution for In Situ Applications**
Matthieu Dreher and Thomas Peterka *(Argonne National Laboratory)*

NVIDIA Industry Session - HPC for AI Development
Time: 16:00-17:00
Room: Grand Hall A
Chair: Jerry Chou *(National Tsinghua University)*

16:00-17:00 **HPC for AI Development**
Simon See *(NVIDIA)*

Panel Session - HPC vs. Big Data: Different Worlds or Common Ground?
Time: 17:00-18:00
Room: Grand Hall
Moderator: Mary Hall *(University of Utah)*

The platforms and programming systems used in high-performance computing and big data analytics are significantly different, even though issues of scalability are fundamental to both areas. In this panel, we will highlight the differences in how these fields have currently evolved, looking specifically at computer architecture, programming systems, and applications. We will also look to the future in understanding how such differences may diminish over time, as similarities potentially outweigh differences.

Panelists: Gabriel Antoniu *(INRIA)*, Todd Gamblin *(Lawrence Livermore National Laboratory)*, Satoshi Matsuoka *(Tokyo Technical University)*, and Michela Taufer *(University of Delaware)*

Gala Dinner
Time: 19:00-21:00
Room: Grand Hall

September 15, 2016 (Thu.)

Keynote 3

Time: 09:00-10:00
Room: Grand Hall
Chair: Chung-Ta King *(National Tsinghua University)*

09:00-10:00 **Who is Afraid of I/O? Exploring I/O Challenges and Opportunities at the Exascale**
Michela Taufer *(University of Delaware)*

Coffee Break
Time: 10:00-10:30

Paper Session 9 - I/O Optimization

Time: 10:30-12:30
Room: Grand Hall
Chair: Matthieu Dorier *(Argonne National Laboratory)*

10:30-11:00 **Exploring Data Migration for Future Deep-Memory Many-Core Systems**
Swann Perarnau and Judicael Zounmevo *(Argonne National Laboratory),* Balazs Gerofi *(RIKEN AICS),* and Kamil Iskra and Pete Beckman *(Argonne National Laboratory)*

11:00-11:30 **GraphMeta: A Graph-based Engine for Managing Large-Scale HPC Rich Metadata**
Dong Dai and Yong Chen *(Texas Tech University),* Philip Carns and John Jenkins *(Argonne National Laboratory),* Wei Zhang *(Texas Tech University),* and Robert Ross *(Argonne National Laboratory)*

11:30-11:45 **Extending SSD Lifetime with Persistent In-memory Metadata Management**
Qingsong Wei, Cheng Chen, Mingdi Xue, Chundong Wang, and Jun Yang *(Data Storage Institute)*

11:45-12:00 **FlashStager: Improving the Performance of SSD-based Data Staging Systems via Write Redirection**
Xuechen Zhang *(Washington State University Vancouver)*, Fang Zheng *(IBM Research)*, and Karsten Schwan *(Georgia Institute of Technology)*

12:00-12:30 **Performance Optimization for All Flash Scale-out Storage**
Myoungwon Oh, Jungyeon Yoon, and Jugwan Eom *(SKTelecom)*; Heon Y. Yeom *(Seoul National University)*; and Jae Yeun Yun and Seungmin Kim (SKTelecom)

Paper Session 10 - Parallel Applications

Time: 10:30-12:30
Room: Grand Hall A
Chair: Naoya Maruyama *(RIKEN)*

10:30-11:00 **Distributed Parallel #SAT Solving**
Jan Burchard, Tobias Schubert, and Bernd Becker *(Albert Ludwigs University Freiburg)*

11:00-11:30 **Smart-MLlib: A High-Performance Machine-Learning Library**
David Siegal, Jia Guo, and Gagan Agrawal *(Ohio State University)*

11:30-11:45 **Distributed Bayesian Probabilistic Matrix Factorization**
Tom Vander Aa and Imen Chakroun *(imec)* and Tom Haber *(UHasselt)*

11:45-12:00 **A Software-Defined Storage for Workflow Applications**
Samer Al-Kiswany *(University of Waterloo)* and Matei Ripeanu *(University of British Columbia)*

12:00-12:30 **VOLAP: A Scalable Distributed System for Real-Time OLAP with High Velocity Data**
Frank Dehne and David Robillard *(Carleton University)* and Andrew Rau-Chaplin and Neil Burke *(Dalhousie University)*

Lunch Break
Time: 12:30-14:00

Student Mentoring Program

Time: 12:30-14:00
Room: Grand Hall A

Paper Session 11 - Data Analytics

Time: 14:00-15:30
Room: Grand Hall
Chair: Matthieu Dreher *(Argonne National Laboratory)*

14:00-14:30 **Horme: Random Access Big Data Analytics**
Guangchen Ruan and Beth Plale *(Indiana University)*

14:30-15:00 **In-Cache MapReduce: Leverage Tiling to Boost Temporal Locality-Sensitive MapReduce Computations**
Daniel Magro and Herve Paulino *(NOVA University of Lisbon)*

15:00-15:15 **Intra-host Rate Control with Centralized Approach**
Zhuang Wang, Ke Liu, Yifan Shen, and Long Li *(Institute of Computing Technology, Chinese Academy of Sciences)*; Jack Y. B. Lee *(The Chinese University of Hong Kong)*; and Mingyu Chen and Lixin Zhang *(Institute of Computing Technology, Chinese Academy of Sciences)*

15:15-15:30 **Fast Big Data Analysis in Geo-Distributed Cloud**
Yue Li and Laiping Zhao *(Tianjin University)*; Chenzhou Cui *(National Astronomical Observatories, CAS (NAOC))*; and Ce Yu *(Tianjin University)*

Paper Session 12 - Optimized Checkpointing

Time: 14:00-15:30
Room: Grand Hall A
Chair: Torsten Hoefler *(ETH Zurich)*

14:00-14:30 **A Lightweight Causal Message Logging Protocol to Lower Fault Tolerance Overhead**
Jinmin Yang *(Hunan University)*

14:30-15:00 **Design and Implementation for Checkpointing of Distributed Resources using Process-level Virtualization**
Kapil Arya *(Mesosphere, Inc.)*; Rohan Garg (Northeastern University); Artem Y. Polyakov *(Mellanox.com)*; and Gene Cooperman *(Northeastern University)*

15:00-15:30 **Deduplication Potential of HPC Applications' Checkpoints**
Jürgen Kaiser, Ramy Gad, Tim Süß, Federico Padua, Lars Nagel, and André Brinkmann *(Johannes Gutenberg University Mainz)*

Coffee Break
Time: 15:30-16:00

Paper Session 13 - Big Data

Time: 16:00-18:00
Room: Grand Hall
Chair: Tim Suess *(Johannes Gutenberg University)*

16:00-16:30 **A Comparative Survey of the HPC and Big Data Paradigms: Analysis and Experiments**
HamidReza Asaadi, Dounia Khaldi, and Barbara Chapman *(Stony Brook University)*

16:30-17:00 **Spark versus Flink: Understanding Performance in Big Data Analytics Frameworks**
Ovidiu-Cristian Marcu *(INRIA Rennes - Bretagne Atlantique)*, Alexandru Costan *(IRISA/INSA Rennes)*, Gabriel Antoniu *(INRIA Rennes - Bretagne Atlantique)*, and Maria S. Perez-Hernandez *(Universidad Politecnica de Madrid)*

17:00-17:15 **Results of a Model for Hadoop YARN MapReduce Task**
Thomas Bressoud and Qiuyi Tang *(Denison University)*

17:15-17:30 **Application-assisted Writeback for Hadoop Clusters**
Jungi Jeong *(KAIST)*, Daewoo Lee *(National Security Research Institute)*, and Seungryoul Maeng *(KAIST)*

17:30-18:00 **Custody: Towards Data-Aware Resource Sharing in Cloud-Based Big Data Processing**
Shiyao MA, Jingjie JIANG, and Bo LI *(Hong Kong University of Science and Technology)* and Baochun LI *(University of Toronto)*

Paper Session 14 - Threads: Heterogeneity and Tasking

Time: 16:00-18:00
Room: Grand Hall A
Chair: Felix Wolf *(TU Darmstadt)*

16:00-16:30 **ARCS: Adaptive Runtime Configuration Selection for Power-Constrained OpenMP Applications**
Md Abdullah Shahneous Bari *(University of Houston),* Nicholas Chaimov *(University of Oregon),* Abid M. Malik *(University of Houston),* Kevin A. Huck *(University of Oregon),* Barbara Chapman *(University of Houston),* Allen Malony *(University of Oregon),* and Osman Sarood *(Yelp Inc.)*

16:30-17:00 **A Review of Lightweight Thread Approaches for High Performance Computing**
Adrián Castelló *(Universitat Jaume I)*, Antonio J. Peña *(Barcelona Supercomputing Center)*, Sangmin Seo *(Argonne National Laboratory)*, Rafael Mayo (Universitat Jaume I), Pavan Balaji *(Argonne National Laboratory)*, and Enrique S. Quintana-Ortí *(Universitat Jaume I)*

17:00-17:15 **Directive-Based Pipelining Extension for OpenMP**
Xuewen Cui *(Virginia Tech)*, Thomas R. W. Scogland and Bronis R. de Supinski *(Lawrence Livermore National Laboratory)*, and Wu-chun Feng *(Virginia Tech)*

17:15-17:30 **Serving More GPU Jobs, with Low Penalty, using Remote GPU Execution and Migration**
Pak Markthub, Akihiro Nomura, and Satoshi Matsuoka *(Tokyo Institute of Technology)*

17:30-17:45 **VarySched: A Framework for Variable Scheduling in Heterogeneous Environments**
Tim Suess, Nils Doering, Ramy Gad, Lars Nagel, and André Brinkmann *(JGU Mainz)*; Dustin Feld and Thomas Soddemann *(Fraunhofer SCAI)*; and Stefan Lankes *(RWTH Aachen)*

17:45-18:00 **Extending the Roofline Model for Asynchronous Many-Task Runtimes**
Joshua Suetterlein *(UDEL)*; Joshua Landwehr, Andres Marquez, and Joseph Manzano *(PNNL)*; and Guang Gao *(UDEL)*

Workshop

Workshop Keynote – It's Not My Fault! Finding Errors in Parallel Codes

David Abramson ***(University of Queensland)***

Abstract

Debugging software has always been difficult, with little tool support available. Finding faults in parallel programs is even harder because the machines and problems are so large, and the amount of state to be examined becomes prohibitive. Faults are often introduced when codes are modified, the software or hardware environment changes or they are scaled up to solve larger problems. All too often we hear the programmers scream "It's not my fault!"

Over the years we have developed a technique called "Relative Debugging", in which a code is debugged against another, reference, version. This makes the process simpler because programmers can compare the state of computation between a faulty version and a previous code that is correct, and the programmer doesn't need to have a mental model of what the program state should be. However, relative debugging can also be expensive because it needs to compare large data structures across the machine. Parallel computers offer a way of accelerating the comparisons using parallel algorithms, making the technique practical.

In this talk I will introduce relative debugging, show how it assists test and debug, and discuss the various techniques used to scale it up to very large problems and machines.

Biography

Director, Research Computing Centre David has been involved in computer architecture and high performance computing research since 1979. He has held appointments at Griffith University, CSIRO, RMIT and Monash University. Prior to joining UQ, he was the Director of the Monash e-Education Centre, Science Director of the Monash e-Research Centre, and a Professor of Computer Science in the Faculty of Information Technology at Monash. From 2007 to 2011 he was an Australian Research Council Professorial Fellow. David has expertise in High Performance Computing, distributed and parallel computing, computer architecture and software engineering. He has produced in excess of 200 research publications, and some of his work has also been integrated in commercial products. One of these, Nimrod, has been used widely in research and academia globally, and is also available as a commercial product, called EnFuzion, from Axceleon. His world-leading work in parallel debugging is sold and marketed by Cray Inc, one of the world's leading supercomputing vendors, as a product called ccdb. David is a Fellow of the Association for Computing Machinery (ACM), the Institute of Electrical and Electronic Engineers (IEEE), the Australian Academy of Technology and Engineering (ATSE), and the Australian Computer Society (ACS). He is currently a visiting Professor in the Oxford e-Research Centre at the University of Oxford.

FTS 2016 –The Second International Workshop on Fault Tolerant Systems

September 15, 2016 (Thu.)

Keynote Speech

Time: 14:00-15:00
Room: Grand Hall B
Chair: Vaidy Sunderam *(Emory University)*

14:00-15:00 **It's Not My Fault! Finding Errors in Parallel Codes**
David Abramson *(University of Queensland)*

FTS 2016 – The Second International Workshop on Fault Tolerant Systems

Time: 15:00-18:00
Room: Grand Hall B
Chair: Vaidy Sunderam *(Emory University)*

15:00-15:30 **Selective Replication for Fault-tolerant Task-Parallel HPC Applications**
Omer Subasi, Gulay Yalcin, Ferad Zyulkyarov, Osman Unsal and Jesus Labarta

Coffee Break
Time: 15:30-16:00

16:00-16:30 **TwinPCG: Dual Thread Redundancy with Forward Recovery for Preconditioned Conjugawte Gradient Methods**
Kiril Dichev and Dimitrios Nikolopoulos

16:30-17:00 **An ABFT Scheme Based on Communication Characteristics**
Upama Kabir and Dhrubajyoti Goswami

17:00-17:30 **Separation Kernel Robustness Testing: The XtratuM Case Study**
Stephen Grixti, Nicholas Sammut, Maria Hernek, Elena Carrascosa, Miguel Masmano and Alfons Crespo

17:30-17:55 **Adaptive Impact-Driven Detection of Silent Data Corruption for HPC Applications**
Sheng Di

17:55-18:00 **Wrap Up**

Conference Information

Conference Venue

Palais de Chine Hotel

Address: No.3, Section 1, Chengde Road, Taipei 103, Taiwan

Tel.: +886-2-2181-9999

Registration

Location:	**5F, Palais de Chine Hotel**	
Operating Hours:	September 13	08:00-17:00
	September 14	08:30-17:00
	September 15	08:30-15:00

Conference Posters

Poster Setup:	September 13	15:00-18:00
Poster Session:	September 13	18:30-20:00
	September 14	09:00-16:00
	September 15	09:00-11:00
Poster Removal:	September 15	11:00-12:00

Internet Access

Free WI-FI is available throughout the conference venue provided by Palais de Chine Hotel.
SSID: palaisdechine (no password required)

Lunch

Conference lunch will be provided for preregistered Student Program attendees in the meeting room.

Palais de Chine Hotel has three restaurants located on the 6th and 17th floor. Participants may also visit a large selection of restaurants at the following venues:

• **Q Sqaure Mall**
Business Hours: 11:00-21:30
Restaurant Locations: 1F & 4F
Distance from Palais de Chine Hotel: Same building
Address: No. 1, Sec. 1, Chengde Road

• **Breeze Taipei Station**
Business Hours: 11:00-21:30
Restaurant Locations: 1F & 4F
Distance from Palais de Chine Hotel: Across the street
Address: No. 3, Be-Ping West Road

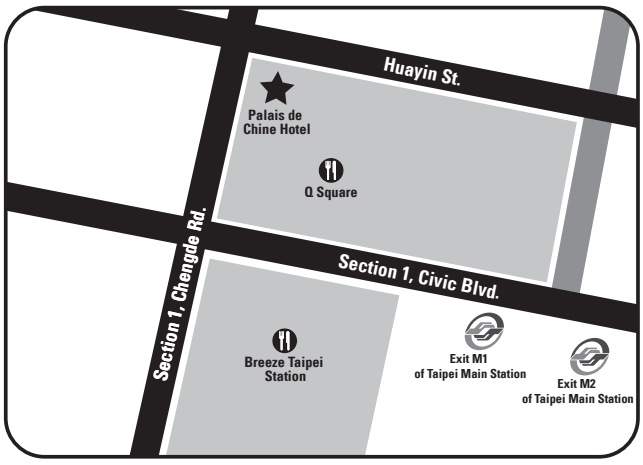

Welcome Reception

Date: Tuesday, September 13
Time: 18:30-20:00
Room: Grand Hall+Grand Hall B
Note: Coupon is required for admission

Gala Dinner

Date: Wednesday, September 14
Time: 19:00-21:00
Room: Grand Hall
Note: Coupon is required for admission

Floor Plan

5F

Palais de Chine Hotel

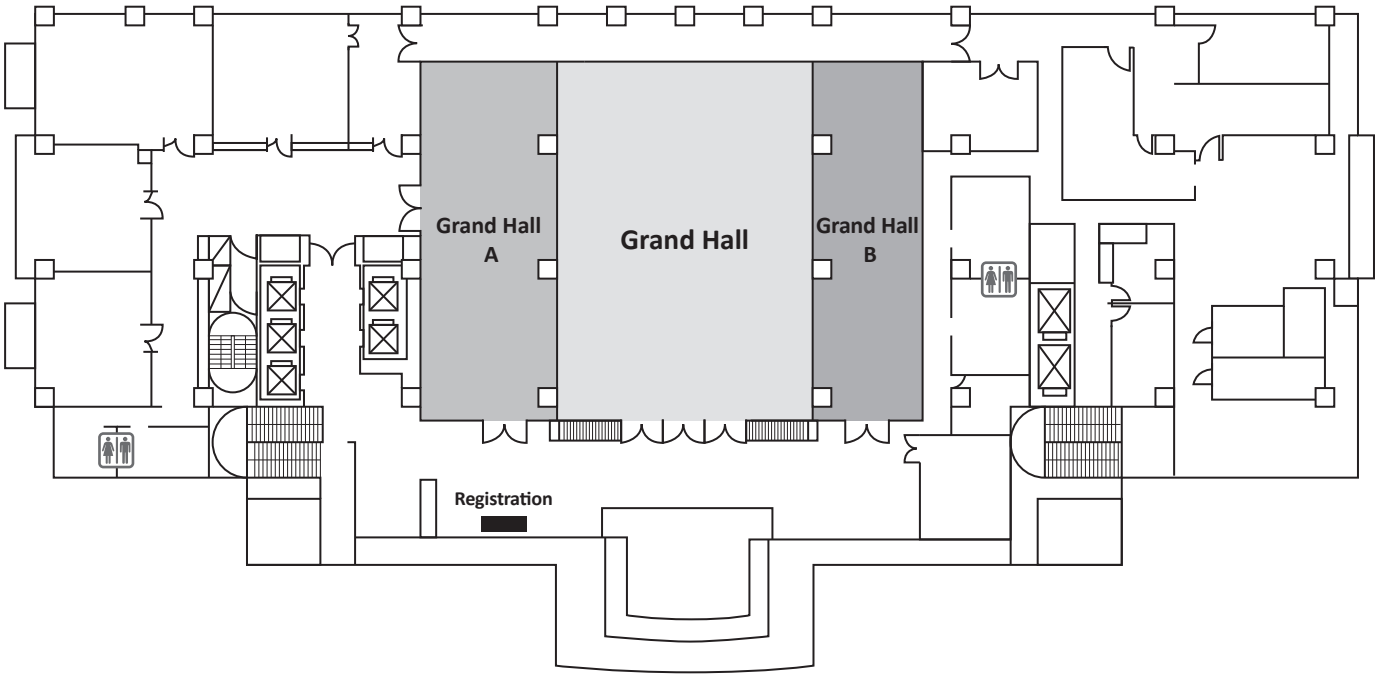

Grand Hall
Opening Remarks, Keynote, Best Paper Award Session, Paper Session, Panel Session, Gala Dinner

Grand Hall+Grand Hall B
Welcome Reception

Grand Hall A
Paper Session, Student Mentoring Program, Tutorials, Industry Session

Grand Hall B
Conference Posters, Workshop