

Table of Contents

Welcome Message from Chairs.....	
Program Agenda.....	
Keynote Speeches.....	
Panel Sessions.....	
Technical Program.....	
Workshop Program.....	
Conference Venue.....	
Transportation Information.....	
Social Programs.....	
Conference Lunch.....	
Sponsors.....	

Dear ICPP 2011 Attendees,

Welcome to the 40th anniversary of the International Conference on Parallel Processing. The ICPP series of conferences has established its reputation and prestige over its long history and has served as one of the most important platforms for researchers and developers to exchange innovative ideas and significant research results in the fields of computer architecture and parallel computing. The conference also provides unique opportunities for researchers, engineers, practitioners, and developers to explore new ideas, techniques, and tools.

The 40th ICPP conference, held during September 13-16, 2011, is co-organized by National Tsing Hua University and National Chiao Tung University, Taiwan, and many volunteers around the world. The conference organization consists of 11 tracks, each handled by one or two Program Vice-Chairs. We have received 363 valid submissions, each of which was thoroughly reviewed by 3-5 reviewers. We are very thankful to hundreds of TPC members who helped out these huge review works. The final technical program consists of very selective 81 papers from 363 submissions.

This year's program includes three outstanding keynote speeches, technical sessions, and panel discussions. There are also 6 workshops in conjunction with ICPP 2011. We would like to thank our organizing committee members, TPC members, reviewers, workshop organizers, workshop TPC members, award chair, publications chair, publicity chairs, financial chair, local arrangement chairs, international liaison chairs, registration chairs, and steering committee chair for their contributions.

We would also like to take this opportunity to thank our sponsors for their strong supports to make this event even more successful: National Chiao Tung University, National Tsing Hua University, Chung Yuan Christian University, Department of Information and Tourism, Chunghwa Telecom, National Science Council, Bureau of Foreign Trade, Ministry of Education, Taiwan Association of Cloud Computing, National Center for High Performance Computing, Industrial Technique Research Institute, Microsoft Research Asia, Chung Yuan Christian University, Institute for Information Industry, Institute of Information Science and Research Center for Information Technology Innovation of Academia Sinica, and NCTU/ITRI Joint Research Center.

In addition to technical programs, we like to urge you to participate in the social programs that we have arranged. Besides, take your time to experience Taiwan's foods, cultures, landscapes, etc. We truly hope you have a memorable experience.

Sincerely,

General Co-Chairs
Wen-Tsuen Chen, National Tsing Hua University
Xiaodong Zhang, Ohio State University

Program Co-Chairs
Guang R. Gao, University of Delaware
Yu-Chee Tseng, National Chiao Tung University

Program Agenda

September 13, 2011 (Tuesday)					
ROOM	B2-1	B2-2	B2-3	B2-4	
08:00-09:00	Registration				
09:00-12:00	W1: CloudSec	W3: P2S2 (08:45-12:20)	W4: EMS	W6: AWASN	
12:00-13:30	Lunch				
13:30-17:00	W2: SRMPDS	W3: P2S2	W5: PSTI	W6: AWASN	
September 14, 2011 (Wednesday)					
ROOM	B2-Howard Hall	B2-1	B2-2	B2-3	B2-4
08:20-08:50	Registration				
08:50-09:00	Opening				
09:00-10:00	Keynote 1				
10:00-10:30	Coffee Break				
10:30-12:00	Panel 1 (10:30-11:30) ----- Panel 2 (11:30-12:30)	S1: Architecture	S2: Wireless		
12:00-13:30	Lunch				
13:30-15:00		S3: Architecture	S4: Wireless	S5: Performance and Modeling	S6: Compilers
15:00-15:30	Coffee Break				
15:30-17:00		S7: Architecture	S8: Wireless	S9: Performance and Modeling	S10: Compilers
18:00-20:00	Reception (Location: Rainbow Terrace the Atrium, 1F)				
September 15, 2011 (Thursday)					
ROOM	B2-Howard Hall	B2-1	B2-2	B2-3	B2-4
08:30-09:00	Registration				
09:00-10:00	Keynote 2				
10:00-10:30	Coffee Break				
10:30-12:30	S11: Cloud Computing	S12: Cluster and Grid Computing	S13: Algorithms	S14: Multi-core and Parallel Systems	
12:30-13:30	Lunch				
13:30-15:00		S15: Cluster and Grid Computing	S16: Algorithms	S17: Multi-core and Parallel Systems	S18: Mobile Computing and Networks (13:30-16:00)
15:00-15:30	Coffee Break				
15:30-17:00		S19: Cluster and Grid Computing	S20: Algorithms	S21: Multi-core and Parallel Systems	
18:00-20:30	Banquet (Location: Howard Hall, B2)				
September 16, 2011 (Friday)					
ROOM	B2-Howard Hall	B2-1	B2-2	B2-3	
08:30-09:00	Registration				
09:00-10:00	Keynote 3				
10:00-10:30	Coffee Break				
10:30-12:30	S22: Cloud Computing	S23: OS and Runtime Technology	S24: Algorithms	S25: P2P Computing and Services	

Keynote 1

The Cloud, the Client and Big Data



Time: 09:00-10:00, September 14 (Wednesday)

Room: B2-Howard Hall

Speaker: Dr. Dennis Gannon

Director of Cloud Research Engagements

eXtreme Computing Group

Microsoft Research

Dr. Dennis Gannon is Director of Cloud Research Engagements for the Microsoft Technical Policy Group. Dr. Gannon's research interests include cloud computing, data analytics and “big data” platforms, large-scale cyberinfrastructure, distributed computing, parallel programming, computational science and problem solving environments. At Microsoft he and his team are working with the research community to demonstrate the potential of cloud computing to enable broad access to data-intensive scientific research.

Prior to coming to Microsoft, Dr. Gannon was a professor and former chair of computer science at Indiana University and the Science Director for the Indiana Pervasive Technology Labs. He has published over 100 refereed articles and he has co-edited 3 books. Dr. Gannon received his Ph.D. in Computer Science from the University of Illinois Urbana-Champaign after receiving a Ph.D. in Mathematics from the University of California, Davis.

Abstract

The first paradigm of science was experimental. This was quickly followed by the second paradigm, theory, to explain the results of experiments. The third paradigm was computation which allows us to explore theory where experimentation is difficult or impossible. There is a fourth paradigm that can be described as deriving new knowledge from massive amounts of data even in cases where we may have very little theory to guide us. This is important because almost every branch of academic research is inundated by the data deluge and basic research methods have to evolve rapidly to cope with it. Access to massive amounts of digital data has already transformed the IT industry. Massive scale data clouds designed to index the web have transformed the advertising and publishing industry. We have mobile client devices that have applications that give us total information about where we are at any given instant including where to eat and where to catch a cab. Our computers are

learning to see and recognize us as we walk through instrumented spaces. That capability is the result of machine learning applied to massive data collections. However, the academic research community is lagging behind in this revolution. While the most adventurous researchers have access to massive supercomputing facilities and communities like high energy physics have well established data analysis pipelines, the majority of researchers limit the scope of their research to what they can do with the computer on their desk. In this talk we will discuss an approach to removing this limitation by building cloud-based data analytics services that are easy to use from the researchers desktop.

Keynote 2

Execution Models without Borders



Time: 09:00-10:00, September 15 (Thursday)

Room: B2-Howard Hall

Speaker: Dr. Thomas Sterling

School of Informatics and Computing and Pervasive

Technology Institute Indiana University

Fellow, Computer Science Research Institute Sandia National Laboratories

Dr. Thomas Sterling holds the position of Professor of Informatics and Computing at the Indiana University (IU) School of Informatics and Computing as well as serving as Director of the Laboratory for System Science and Engineering at the IU Center for Research in Extreme Scale Technology (CREST). He also is an Adjunct Professor at the Louisiana State University (LSU) Center for Computation and Technology (CCT) and CSRI Fellow at Sandia National Laboratories. Since receiving his Ph.D from MIT in 1984 as a Hertz Fellow he has engaged in applied research in related fields associated with parallel computing system structures, semantics, and operation in industry, government labs, and academia. Dr. Sterling is best known as the "father of Beowulf" for his pioneering research in commodity/Linux cluster computing. He was awarded the Gordon Bell Prize in 1997 with his collaborators for this work. Thomas Sterling currently leads the ParalleX Research Group to devise a new model of computation establishing the foundation principles to guide the co-design for the development of future generation Exascale computing systems by the end of this decade. His research has been sponsored by NSF, NASA, NSA, DOE, DARPA, Army Corps of Engineers, and Microsoft. He is the co-author of six books and holds six patents.

Abstract

Technology trends have forced parallel processing to new architecture structures that stress conventional usage practices beyond effective applications. The bounds on increased clock rates due to power constraints and increased processor complexity due to ILP limitations have forced multi/manycore structures in combination with GPU accelerators as typified by three of the top four MPP systems in the world. Once the Communicating Sequential Process (CSP) execution model is reflected by the

MPI programming interface dominated both MPP and commodity cluster systems, now the parallel processing community struggles to find alternative methods to effectively program and manage such heterogeneous parallel systems. Historically, the field of HPC has experienced 5 previous phase changes where technology and the need for new classes or architecture and programming methods to exploit it, the most recent more than two decades ago. The fundamental issue that challenges the international parallel processing community is the new HPC phase change that will push us into the era of nano-scale technology and multi-billion-way parallelism by the end of this decade. Such a paradigm shift is critical and its lack of adoption is already impeding progress in hardware and software system codesign and extreme scale application development. The critical performance factors that must be addressed are the interrelated behavior properties of starvation, latency, overhead, and contention. As in the past, any new execution model must mitigate their effects of performance degradation. Also, like the past vector model, SIMD, and CSP execution models, the future model must be international in its adoption, general in its breadth of application, and a powerful tool across industrial manufacturers and ISV suppliers. Essentially all borders must be crossed by the future generation execution model. This Keynote Address will focus on this challenge that affects every nation, every user community, every computational challenge, and every producer. The ParalleX experimental execution model will be described as a concrete exemplar of one possible new paradigm that serves as a framework to synthesize prior art while incorporating innovation in order to provide a way forward, eliminating borders and empowering future cooperation among peoples, domains, disciplines, and standards.

Keynote 3

The "Single-chip Cloud Computer", an IA Tera-scale Research Processor



Time: 09:00-10:00, September 16 (Friday)

Room: B2-Howard Hall

Speaker: Dr. Jim Held

Intel Fellow

Director, Tera-Scale Computing Research Intel Labs

Jim Held is an Intel Fellow who leads a virtual team of architects conducting Tera-Scale Computing Research in Intel's Labs. Since joining Intel in 1990, he has led research and development in a variety of Intel's architecture labs concerned with media and interconnect technology, systems software, multi-core processor architecture and virtualization. Before coming to Intel, Jim worked in research and teaching capacities in the Medical School and Department of Computer Science at the University of Minnesota where he earned a Ph.D. (1988) in Computer and Information Science.

Abstract

Intel Labs has created a second generation experimental "Single-chip Cloud Computer," (SCC) that contains the most Intel Architecture cores ever integrated on a silicon CPU chip - 48 cores. SCC is a concept vehicle, incorporating technologies intended to scale multi-core processors to 100 cores and beyond, such as an on-chip network, advanced power management technologies and support for message-passing. Architecturally, SCC is a microcosm of a cloud datacenter. Each core can run a separate OS and software stack and act like an individual compute node that communicates with other compute nodes over the on-die network fabric, thus supporting the "scale-out" message-passing programming models that have been proven to scale to 1000s of processors.

The SCC also serves as an experimental platform for a wide range of parallel computing research on scalable programming models and architectures and increasing our understanding of how to build better processors for the Cloud. It is currently being used by a worldwide community of academic and industry co-travelers.

This talk will describe the architecture of the SCC platform and discuss its role in the broader context of our Tera-scale research.

Panel 1

Programming Environments at Extreme Scale

Time: 10:30-11:30, September 14 (Wednesday)

Room: B2-Howard Hall

Exa-scale systems are expected to present an extremely large challenge to computational scientists attempting to leverage the full extent of the capabilities such systems will provide. The systems are expected to have on the order of ten million computational elements, with order several-hundred such cores per node, include multiple types of computational elements, and have non-uniform memory access characteristics with deep memory hierarchies. Applications are also expected to need to use on the order of ten to one-hundred way concurrency per core, provided by fine-grain multithreading. The challenges for application developers trying to use such systems is to understand in detail how their codes perform on such systems, and how these codes need to be changed to make better use of such systems. This panel will discuss current and future research directions aimed at helping application developers develop applications for such systems, understand the performance of such applications on these systems, and help in transforming these applications to better utilize such systems.

Panel Coordinator:

Richard L. Graham, Oak Ridge National Laboratory

Panelists:

Barbara Chapman, University of Houston

Robert Harrison, Oak Ridge National Laboratory

Wolfgang Nagel, The Technical University of Dresden

Martin Schulz, Lawrence Livermore National Laboratory

Panel 2

Opportunities and Challenges in Multi-core Era: A Cross-Layer Dialogue

Time: 11:30-12:30, September 14 (Wednesday)

Room: B2-Howard Hall

Processors have been moved to multi-core architectures for many years. However, it is still very challenging today to build a scalable, high performance, low-power, and cost-effective multi-core system. The shift towards multi-core not only poses challenges for computer architectures, but also brings new issues distributed across several communities such as applications and algorithms, programming models, languages and compilers, virtual machines, operating systems and run-time supports. This panel gives us an opportunity to stimulate discussions and dialogue across the spectrum of different research communities. We would like to invite researchers from parallel algorithms and applications, programming languages and compiler designs, OSs and run-time supports, and computer architectures to share their ideas from different angle of view, and join our discussion on challenges and research directions of multi-core systems.

Panel Coordinator:

Shiao-Li Tsao, National Chiao Tung University

Panelists:

Tien-Fu Chen, National Chiao Tung University

Jürg Gutknecht, ETH Zurich

Jim Held, Intel Labs

Chung-Ta King, National Tsing Hua University

Jane Win-Shih Liu, Academia Sinica

S1: Architecture

Time: 10:30-12:00, September14 (Wednesday)

Room: B2-1

Session Chair: Prof. Krishna Kavi, University of North Texas

S1-1 A DFA with Extended Character-set for Fast Deep Packet Inspection

Cong Liu, Ai Chen, Di Wu, Jun Zhang and Jie Wu

S1-2 Symbiotic Scheduling for Shared Caches in Multi-Core Systems Using Memory Footprint Signature

Mrinmoy Ghosh, Ripal Nathuji, Min Lee, Karsten Schwan and Hsien-Hsin Lee

S1-3 A Distributed Switch Architecture for On-Chip Networks

Antoni Roca, Carles hernández, José Flich, Federico Silla and José Duato

S2: Wireless/Sensor Networks and Pervasive Computing

Time: 10:30-12:00, September14 (Wednesday)

Room: B2-2

Session Chair: Prof. Meng-Shiuan Pan, Tamkang University

S2-1 Patrolling Mechanisms for Disconnected Targets in Wireless Mobile Data Mules Networks

Chih-Yung Chang, Chih-Yu Lin, Cehn-Yu Hsieh and Yi-Jung Ho

S2-2 Peer-to-Peer Object Tracking in the Internet of Things

Yanbo Wu, Quan Z. Sheng and Damith Ranasinghe

S2-3 An Innovative Scheme for Increasing Connectivity in ZigBee Networks

Chia-Ming Wu, Ruay-Shiung Chang and Pu-I Lee

S3: Architecture

Time: 13:30-15:00, September14 (Wednesday)

Room: B2-1

Session Chair: Prof. Chung-Ping Chung, National Chiao Tung University

S3-1 Evaluation of Techniques to Improve Cache Access Uniformities

Izuchukwu Nwachukwu, Kavi Krishna, Ademola Fawibe and Chris Yan

S3-2 Energy- and Performance-Efficient Thread Mapping in NoC-based CMPs under Process Variations

Carles Hernández, Federico Silla and Jose Duato

S3-3 Energy-Efficient Cache Coherence Protocols in Chip-Multiprocessors for Server Consolidation

Antonio García-Guirado, Ricardo Fernández-Pascual, Alberto Ros and José M. García

S4: Wireless/Sensor Networks and Pervasive Computing

Time: 13:30-15:00, September14 (Wednesday)

Room: B2-2

Session Chair: Prof. Tomotaka Wada, Kansai University

S4-1 A Distributed Flow-Based Guiding Protocol in Wireless Sensor Networks

Po-Yu Chen, Zan-Feng Kao, Wen-Tsuen Chen and Chi-Han Lin

S4-2 Efficient Bandwidth Allocation with QoS Guarantee for IEEE 802.16 Systems

Da-Nung Lai, Tsung-Chuan Huang and Hung-Yi Chi

S4-3 Gradient-based Aggregation in Forest of Sensors (GrAFS)

Ravi Prakash and Ehsan Nourbakhsh

S5: Performance and Modeling

Time: 13:30-15:00, September14 (Wednesday)

Room: B2-3

Session Chair: Dr. Martin Schulz, Lawrence Livermore National Laboratory

S5-1 Unveiling Internal Evolution of Parallel Application Computation Phases

Harald Servat, Germán Llort, Judit Giménez, Kevin Huck and Jesús Labarta

S5-2 Cache Pirating: Measuring the Curse of the Shared Cache

David Eklov, Nikos Nikoleris, David Black-Schaffer and Erik Hagersten

S5-3 Parallel Performance Measurement of Heterogeneous Parallel Systems with GPUs

Allen Malony, Scott Biersdorff, Sameer Shende, Heike Jagode, Stanimire Tomov, Guido Juckeland, Robert Dietrich, Duncan Poole and Christopher Lamb

S6: Compilers, Programming Models and Languages

Time: 13:30-15:00, September14 (Wednesday)

Room: B2-4

Session Chair: Dr. Brice Goglin, INRIA

S6-1 GStream: A General-Purpose Data Streaming Framework on GPU Clusters

Yongpeng Zhang and Frank Mueller

S6-2 CAB: Cache Aware Bi-tier Task-stealing in Multi-socket Multi-core Architecture

Quan Chen, Zhiyi Huang, Minyi Guo and Jingyu Zhou

S6-3 Virtual Topologies for Scalable Resource Management and Contention Attenuation in a Global Address Space Model on the Cray XT5

Weikuan Yu, Vinod Tipparaju, Xinyu Que and Jeffrey Vetter

S7: Architecture

Time: 15:30-17:00, September14 (Wednesday)

Room: B2-1

Session Chair: Prof. Kuo-Wei Hsu, National Chengchi University

S7-1 PEPCP: A Power-Efficient Parallel Coherence Protocol for Large-Scale Network-on-Chip

Fucen Zeng, Lin Qiao and Wei Wang

S7-2 Eager Meets Lazy: The Impact of Write-Buffering on Hardware Transactional Memory

Anurag Negi, Ruben Titos-Gil, Manuel E. Acacio, Jose M. Garcia and Per Stenstrom

S7-3 Tolerating Load Miss Latency by Extending Effective Instruction Window with Low Complexity

Walter Li, Chin-Ling Huang and Chung-Ping Chung

S8: Wireless/Sensor Networks and Pervasive Computing

Time: 15:30-17:00, September14 (Wednesday)

Room: B2-2

Session Chair: Prof. Shikharesh Majumdar, Carleton University

S8-1 Unilateral Wakeup for Mobile Ad Hoc Networks

Shan-Hung Wu, Jang-Ping Sheu and Chung-Ta King

S8-2 A Secure Data Aggregation based Trust Management Approach for Dealing with Untrustworthy Motes in Sensor Network

Sanjay Madria

S8-3 Video-Like Compression for High Efficiency Database Storage of Wireless Sensor Networks

Niang-Ying Huang, Chi-Cheng Chuang and Ray-I Chang

S9: Performance and Modeling

Time: 15:30-17:00, September14 (Wednesday)

Room: B2-3

Session Chair: Prof. Gul Agha, University of Illinois

S9-1 Exposing Complex Bug-Triggering Conditions in Distributed Systems via Graph Mining

Eunsoo Seo, Mohammad Maifi Hasan Khan, Prasant Mohapatra, Jiawei Han and Tarek Abdelzaher

S9-2 Probabilistic Communication and I/O Tracing with Deterministic Replay at Scale

Xing Wu, Karthik Vijayakumar, Frank Mueller, Xiaosong Ma and Philip Roth

S9-3 Interpreting Performance Data Across Intuitive Domains

Martin Schulz, Joshua A. Levine, Peer-Timo Bremer, Todd Gamblin and Valerio Pascucci

S10: Compilers, Programming Models and Languages

Time: 15:30-17:00, September14 (Wednesday)

Room: B2-4

Session Chair: Prof. Kyoung-Woo Lee, Yonsei University

S10-1 A Comprehensive Performance Comparison of CUDA and OpenCL

Jianbin Fang, Ana Lucia Varbanescu and Henk Sips

S10-2 Enabling Multithreading on CGRAs

Jared Payer, Reiley Jeyapaul, Mahdi Hamzeh, Aviral Shrivastava and Sarma Vrudhula

S10-3 Enhancing the Role of Inlining in Effective Interprocedural Parallelization

Jichi Guo, Mike Stiles, Qing Yi and Kleanthis Psarris

S11: Cloud Computing

Time: 10:30-12:30, September15 (Thursday)

Room: B2-Howard Hall

Session Chair: Prof. Yili Gong, Wuhan University

S11-1 Location-aware MapReduce in Virtual Cloud

Yifeng Geng, Shimin Chen, Yongwei Wu, Ryan Wu, Guangwen Yang and Weimin Zheng

S11-2 WAVNet: Wide-Area Network Virtualization Technique for Virtual Private Cloud

Zheming Xu, Sheng Di, Weida Zhang and Cho-Li Wang

S11-3 Virtual Machine Provisioning Based on Analytical Performance and QoS in Cloud Computing Environments

Rodrigo Calheiros, Rajiv Ranjan and Rajkumar Buyya

S11-4 CSR: A Cloud-assisted Speech Recognition Service for Personal Mobile Device

Yu-Shuo Chang, Shih-Hao Hung, Nick Wang and Bor-Shen Lin

S12: Cluster and Grid Computing

Time: 10:30-12:00, September15 (Thursday)

Room: B2-1

Session Chair: Dr. Richard Graham, Oak Ridge National Laboratory

S12-1 IDEA - An API for Parallel Computing with Large Spatial Datasets

Baoqiang Yan and Philip Rhodes

S12-2 Performance of CUDA Virtualized Remote GPUs in High Performance Clusters

José Duato, Antonio J. Peña, Federico Silla, Rafael Mayo and Enrique S. Quintana-Ortí

S12-3 CRFS: A Lightweight User - Level Filesystem for Generic Checkpoint/Restart

Xiangyong Ouyang, Raghunath Rajachandrasekar, Xavier Besseron, Hao Wang, Jian Huang and Dhabaleswar Panda

S13: Algorithms Design and Parallelization

Time: 10:30-12:30, September15 (Thursday)

Room: B2-2

Session Chair: Dr. Hsi-Ya Chang, National Center for High-Performance Computing

S13-1 A Scalable Tridiagonal Solver for GPUs

Hee-Seok Kim, Shengzhao Wu, Li-Wen Chang and Wen-Mei Hwu

S13-2 On the Performance of Greedy Algorithms for Power Consumption Minimization

Anne Benoit, Paul Renaud-Goud and Yves Robert

S13-3 Optimal Data Allocation for Scratch-Pad Memory on Embedded Multi-core Systems

Yibo Guo, Qingfeng Zhuge, Jingtong Hu, Meikang Qiu, Wei-Che Tseng and Edwin H.-M. Sha

S13-4 Energy-aware Mappings of Series-parallel Workflows onto Chip Multiprocessors

Anne Benoit, Paul Renaud-Goud, Yves Robert and Rami Melhem

S14: Multi-core and Parallel Systems

Time: 10:30-12:00, September15 (Thursday)

Room: B2-3

Session Chair: Prof. Taisuke-Arai BOKU, University of Tsukuba

S14-1 Moving Database Systems to Multicore - A Modular Approach

Victor Pankratius and Martin Heneka

S14-2 GSNP: A DNA Single-Nucleotide Polymorphism Detection System with GPU Acceleration

Mian Lu, Jiuxin Zhao, Qiong Luo, Bingqiang Wang, Shaohua Fu and Zhe Lin

S14-3 Understanding Off-chip Memory Contention of Parallel Programs in Multicore Systems

B.M. Tudor, Y.M. Teo and Simon See

S15: Cluster and Grid Computing

Time: 13:30-15:00, September15 (Thursday)

Room: B2-1

Session Chair: Prof. Phillip Dickens, University of Maine

S15-1 Efficient Energy Management using Adaptive Reinforcement Learning-based Scheduling in Large-Scale Distributed Systems

Masnida Hussin, Young Choon Lee and Albert Y. Zomaya

S15-2 QoS Preference-Aware Replica Selection Strategy Using MapReduce-Based PGA in Data Grids

Runqun Xiong, Junzhou Luo, Aibo Song, Bo Liu and Fang Dong

S15-3 Optimizing Process-to-Core Mappings for Two Dimensional Broadcast/Reduce on Multicore Architectures

Christer Karlsson, Teresa Davies, Chong Ding, Hui Liu and Zizhong Chen

S16: Algorithms Design and Parallelization

Time: 13:30-15:00, September15 (Thursday)

Room: B2-2

Session Chair: Prof. Wen-Chih Peng, National Chiao Tung University

S16-1 Modeling and Practical Evaluation of a Service Location Problem in Large Scale Networks

Olivier Beaumont, Nicolas Bonichon and Hubert Larchevêque

S16-2 Optimizing SpMV for Diagonal Sparse Matrices on GPU

Xiangzheng Sun, Yunquan Zhang, Ting Wang, Xianyi Zhang, Liang Yuan and Li Rao

S16-3 Parallel Discovery of Direct Causal Relations and Markov Boundaries with Applications to Gene Networks

Olga Nikolova and Srinivas Aluru

S17: Multi-core and Parallel Systems

Time: 13:30-15:00, September15 (Thursday)

Room: B2-3

Session Chair: Dr. Jan-Jan Wu, Academia Sinica

S17-1 Accelerating Sparse Matrix Vector Multiplication in Iterative Methods Using GPU

Kiran Kumar Matam and Kishore Kothapalli

S17-2 Implications of Merging Phases on Scalability of Multi-core Architectures

Madhavan Manivannan, Ben Juurlink and Per Stenstrom

S17-3 UnSync: A Soft Error Resilient Redundant Multicore Architecture

*Aviral Shrivastava, Reiley Jeyapaul, Fei Hong, Abhishek Rhisheekesan and
Kyoung Lee*

S18: Mobile Computing and Networks

Time: 13:30-16:00, September15 (Thursday)

Room: B2-4

Session Chair: Prof. Sanjay Madria, Missouri University of Science and Technology

S18-1 Understanding the Flooding in Low-Duty-Cycle Wireless Sensor Networks

Zhenjiang Li, Mo Li, Junliang Liu, Yunhao Liu and Shaojie Tang

S18-2 On Using Contact Expectation for Routing in Delay Tolerant Networks

Honglong Chen and Wei Lou

S18-3 Making Many People Happy: Greedy Solutions for Content Distribution

Yunsheng Wang, Yuhong Guo and Jie Wu

S18-4 ALERT: An Anonymous Location-based Efficient Routing Protocol in MANETs

Lianyu Zhao and Haiying Shen

S18-5 Privacy Leakage in Access Mode: Revisiting Private RFID Authentication Protocols

Qingsong Yao, Jinsong Han, Yong Qi, Lei Yang and Yunhao Liu

S19: Cluster and Grid Computing

Time: 15:30-17:00, September15 (Thursday)

Room: B2-1

Session Chair: Prof. Kuan-Chou Lai, National Taichung University

S19-1 An Efficient Programming Paradigm for Shared-Memory Master-Worker Video Decoding on TILE64 Many-Core Platform

*Xuan-Yi Lin, Kuan-Chou Lai, Kuan-Ching Li, Shau-Yin Tseng and Yeh-Ching
Chung*

S19-2 MiF: Mitigating the Intra-file Fragmentation in Parallel File System

Letian Yi, Jiwu Shu, Youyou Lu, Wei Wang and Weimin Zheng

S19-3 Checkpoint and Run-Time Adaptation with Pluggable Parallelisation

Bruno Medeiros and João Sobral

S20: Algorithms Design and Parallelization

Time: 15:30-17:00, September15 (Thursday)

Room: B2-2

Session Chair: Prof. Philip Wilsey, University of Cincinnati

S20-1 Bloom Filter Performance on Graphics Engines

Lin Ma, Roger Chamberlain, Jeremy Buhler and Mark Franklin

S20-2 Kernel Assisted Collective Intra-node MPI Communication Among Multi-core and Many-core CPUs

Teng Ma, George Bosilca, Aurelien Bouteiller, Brice Goglin, Jeffrey Squyres and Jack Dongarra

S20-3 OCL-BodyScan: A Case Study for Application-centric Programming of Many-Core Processors

Ana Lucia Varbanescu, Milos Raskovic, Henk Sips, Maarten Ditzel and Wouter Vlothuizen

S21: Multi-core and Parallel Systems

Time: 15:30-17:30, September15 (Thursday)

Room: B2-3

Session Chair: Prof. Tarek El-Ghazawi, The George Washington University

S21-1 PC-Mesh: A Dynamic Parallel Concentrated Mesh

Jesús Camacho, José Flich, Antoni Roca and José Duato

S21-2 Data-Driven Tasks and their Implementation

Vivek Sarkar and Sagnak Tasirlar

S21-3 Combining Congested-Flow Isolation and Injection Throttling in HPC Interconnection Networks

Jesus Escudero-Sahuquillo, Ernst Gunnar Gran, Pedro Javier Garcia, Jose Flich, Tor Skeie, Olav Lysne, Francisco Jose Quiles and Jose Duato

S21-4 GPU Resource Sharing and Virtualization on High Performance Computing Systems

Teng Li, Vikram Narayana, Esam El-Araby and Tarek El-Ghazawi

S22: Cloud Computing

Time: 10:30-12:30, September 16 (Friday)

Room: B2-Howard Hall

Session Chair: Prof. Jen-Jee Chen, National University of Tainan

S22-1 SQLMR : A Scalable Database Management System for Cloud Computing

Meng-Ju Hsieh, Chao-Rui Chang, Li-Yung Ho, Jan-Jan Wu, Pangfeng Liu and Yeh-Ching Chung

S22-2 S3: An Efficient Shared Scan Scheduler on MapReduce Framework

Lei Shi, Xiaohui Li and Kian-Lee Tan

S22-3 Adaptive Disk I/O scheduling for MapReduce in Virtualized Environment

Shadi Ibrahim, Hai Jin, Lu Lu, Bingsheng He and Song Wu

S22-4 aMOSS: Automated Multi-Objective Server Provisioning with Stress-Strain Curving

Palden Lama and Xiaobo Zhou

S23: OS and Runtime Technology

Time: 10:30-12:00, September 16 (Friday)

Room: B2-1

Session Chair: Prof. Shanq-Jang Ruan, National Taiwan University of Science and Technology

S23-1 LnQ: Building High Performance Dynamic Binary Translators with Existing Compiler Backends

Chun-Chen Hsu, Pangfeng Liu, Chien-Min Wang, Jan-Jan Wu, Ding-Yong Hong, Pen-Chung Yew and Wei-Chung Hsu

S23-2 Memcached Design on High Performance RDMA Capable Interconnects

Jithin Jose, Hari Subramoni, Miao Luo, Minjia Zhang, Jian Huang, Md. Wasi-ur Rahman, Nusrat Islam, Xiangyong Ouyang, Hao Wang, Sayantan Sur and Dhabaleswar Panda

S23-3 Combining Hard Periodic and Soft Aperiodic Real-Time Task Scheduling on Heterogeneous Compute Resources

Hsiang-Kuo Tang, Parmesh Ramanathan and Katherine Compton

S24: Algorithms Design and Parallelization

Time: 10:30-12:00, September 16 (Friday)

Room: B2-2

Session Chair: Prof. Ching-Hsien Hsu, Chung Hua University

S24-1 Memory Mapping and Task Scheduling Techniques for Computation

Models of Image Processing on Many-Core Platform

Ang-Chih Hsieh, Yi-Ta Wu, Shau-Yin Tseng and TingTing Hwang

S24-2 On The Energy Complexity of Parallel Algorithms

Vijay Anand Korthikanti, Gul Agha and Mark Greenstreet

S24-3 Cache Accurate Time Skewing in Iterative Stencil Computations

Robert Strzodka, Mohammed Shaheen, Dawid Pajak and Hans-Peter Seidel

S25: P2P Computing and Services

Time: 10:30-12:00, September 16 (Friday)

Room: B2-3

Session Chair: Dr. Rodrigo Calheiros, University of Melbourne

S25-1 Probabilistic Best-fit Multi-dimensional Range Query in Self-Organizing Cloud

Sheng Di, Cho-Li Wang, Weida Zhang and Luwei Cheng

S25-2 ShareStorm: a High-Performance and ISP-Friendly P2P Content Distribution Protocol

Yingchun Lei, Litang Yang, Yili Gong and Wenjie Wang

S25-3 On the QoS of Offline Download in Retrieving Peer-side File Resource

Yuanjian Xing, Zhi Yang, Chi Chen, Jilong Xue and Yafei Dai

W1: Security in Cloud Computing (CloudSec)

Session I

Time: 09:00-10:00, September13 (Tuesday)

Room: B2-1

Session Chair: Chun-Chieh Huang, Minghsin University of Science and Technology

W1-1 A Generic Scheme for Data Sharing in Cloud

Yanjiang Yang

W1-2 Secure Connectivity for Intra-Cloud and Inter-Cloud Communication

Shiping Chen, Surya Nepal and Ren Liu

Session II

Time: 10:20-12:00, September13 (Tuesday)

Room: B2-1

Session Chair: Chun-Chieh Huang, Minghsin University of Science and Technology

W1-3 A Secure Cloud Backup System with Assured Deletion and Version Control

Arthur Rahumed, Henry C. H. Chen, Yang Tang, Patrick P. C. Lee and John C. S. Lui

W1-4 A Secure File Allocation Algorithm for Heterogeneous Distributed Systems

Yun Tian, Mohammed I. Alghamdi, Jiong Xie, Shu Yin, Ji Zhang, Meikang Qiu, Yiming Yang and Xiao Qin

W1-5 Implications of Recovery Schemes for Virtualization Platform

Guanhua Tian and Dan Meng

W1-6 A Security Framework of Group Location-Based Mobile Applications in Cloud Computing

Yu-Jia Chen and Li-Chun Wang

W2: Scheduling and Resource Management for Parallel and Distributed Systems (SRMPDS)

Session I: GPU and Multi-core Systems

Time: 13:30-14:50, September13 (Tuesday)

Room: B2-1

W2-1 Analyzing the Effects of Multicore Architectures and On-host

Communication Characteristics on Collective Communications

Joshua Ladd, Manjunath Gorentla Venkata, Richard Graham and Pavel Shamis

W2-2 Energy-Aware Workload Consolidation on GPU

Dong Li, Surendra Byna and Srimat Chakradhar

W2-3 The Power Efficiency of GPUs in Multi Nodes Environment with Molecular Dynamics

Takuro Udagawa and Masakazu Sekijima

W2-4 An Efficient I/O Aggregator Assignment Scheme for Collective I/O Considering Processor Affinity

Kwangho Cha and Seungryoul Maeng

Session II: Cloud and Distributed Systems

Time: 15:20-17:00, September13 (Tuesday)

Room: B2-1

W2-5 Agent-based Adaptive Resource Allocation on the Cloud Computing Environment

Gihun Jung and Kwang Mong Sim

W2-6 A Simulation Framework for Reconfigurable Processors in Large-scale Distributed Systems

M. Faisal Nadeem, S. Arash Ostadzadeh, M. Nadeem, J.S.S.M Wong and Klm Bertels

W2-7 An Extensible Design of a Load-Aware Virtual Router Monitor in User Space

Harry F. W. Choi and Patrick P. C. Lee

W2-8 Can MPI Benefit Hadoop and MapReduce Applications?

Xiaoyi Lu, Bing Wang, Li Zha and Zhiwei Xu

W2-9 P2G: A Framework for Distributed Real-Time Processing of Multimedia Data

Håvard Espeland, Paul B. Beskow, Håkon K. Stensland, Preben N. Olsen, Ståle Kristoffersen, Carsten Griwodz and Pål Halvorsen

W3: Parallel Programming Models and Systems Software for High-End Computing (P2S2)

Opening Remarks

Time: 08:45-09:00, September13 (Tuesday)

Room: B2-2

Session Chair: Pavan Balaji, Argonne National Laboratory

Session I: Keynote

Time: 09:00-10:00, September13 (Tuesday)

Room: B2-2

Session Chair: Pavan Balaji, Argonne National Laboratory

W3-1 Exascale: Why It Is Different

Dr. Barbara Chapman, University of Houston

Session II: Programming Models and Runtime Systems

Time: 10:30-12:20, September13 (Tuesday)

Room: B2-2

Session Chair: Martin Schulz, Lawrence Livermore National Laboratory

W3-2 Recomposing An Irregular Algorithm Using a Novel Low-Level PGAS Model

Megan Vance and Peter Kogge

W3-3 A Middleware for Concurrent Programming in MPI Applications

Tobias Berka, Helge Hagenauer and Marian Vajtersic

W3-4 Kangaroo: Reliable execution of Scientific Applications with DAG Programming Model

Kai Zhang, Kang Che and Wei Xue

W3-5 JETS: Language and System Support for Many Parallel Task Computing

Justin Wozniak and Michael Wilde

Session III: Scheduling and Workflows

Time: 13:30-15:20, September13 (Tuesday)

Room: B2-2

Session Chair: Vinod Tipparaju, Oak Ridge National Laboratory

W3-6 Energy-Constrained Dynamic Resource Allocation in a Heterogeneous Computing Environmen

B. Dalton Young, Jonathan Apodaca, Luis Diego Briceno, Jay Smith, Sudeep Pasricha, Anthony A. Maciejewski, Howard Jay Siegel, Bhavesh Khemka, Shirish Bahirat, Adrian Ramirez and Yong Zou

W3-7 Job Co-Scheduling on Coupled High-End Computing Systems

Wei Tang, Narayan Desai, Venkatram Vishwanath, Daniel Buettner and Zhiling Lan

W3-8 Integrating Scientific Workflows and Large Tiled Display Walls: Bridging the Visualization Divide

Hoang Nguyen, David Abramson, Blair Bethwaite, Minh Ngoc Dinh, Colin Enticott, Stephen Firth, Slavisa Garic, Ian Harper, Martin Lackmann, A.B.M. Russel, Stefan Schek and Mary Vail

W3-9 Restricted Admission Control in View-Oriented Transactional Memory

K. Leung and Z. Huang

Session IV: Communication and I/O

Time: 15:40-17:00, September 13 (Tuesday)

Room: B2-2

Session Chair: Brice Goglin, INRIA

W3-10 CellPilot: A Seamless Communication Solution for Hybrid Cell Clusters

Natalie Girard, William Gardner, John Carter and Gary Grewal

W3-11 Interval Based I/O: A New Approach to Providing High Performance Parallel I/O

Jeremy Logan and Phillip Dickens

W3-12 Improving Performance of the Irregular Data Intensive Application with Small Computation Workload for CMPs

Zhimin Gu

W4: Embedded Multicore Systems (EMS)

Session I: Keynote

Time: 09:00-10:00, September 13 (Tuesday)

Room: B2-3

Session Chair: Prof. Shang-Hong Lai

W4-1 Debunking the 100X GPU vs. CPU Myth: An Evaluation of Throughput Computing on CPU and GPU

Dr. Yen-Kuang Chen, Intel

Session II

Time: 10:20-12:00, September 13 (Tuesday)

Room: B2-3

Session Chair: Prof. Pao-Ann Hsiung and Dr. Yanqin Yang

W4-2 Parallelized Face Based RMS System on a Multi-core Embedded Computing Platform

Te-Feng Su, Jia-Jhe Li, Chih-Hsueh Duan, Shu-Fan Wang and Shang-Hong Lai

W4-3 Adaptive Performance Monitoring for Embedded Multicore Systems

Chun-Yi Shih, Ming-Chih Li, Chao-Sheng Lin, Pao-Ann Hsiung, Chih-Hung Chang, William C. Chu, Nien-Lin Hsueh, Chihhsiong Shih, Chao-Tung Yang and Chorng-Shiuh Koong

W4-4 C++ Compiler Supports for Embedded Multicore DSP Systems

Chi-Bang Kuan, Jia-Jhe Li, Chung-Kai Chen and Jenq Kuen Lee

W4-5 Embedded Network Intrusion Detection Systems with a Multi-Core Aware Packet Capture Module

Chia-Hao Hsu and Sheng-De Wang

W4-6 An Efficient Approach of Power Reducing for Scratch-pad Memory based Embedded Systems

Yanqin Yang, Wenchao Xu, Minyi Guo and Zili Shao

W4-7 Accelerating the Near Non-bonded Force Computation in Desmond with Graphic Processing Units, Hualiang Deng

Xin Li, Xiaoguang Liu and Gang Wang

W5: Parallel Software Tools and Tool Infrastructure (PTSI)

Session I

Time: 13:30-14:50, September 13 (Tuesday)

Room: B2-3

Session Chair: Karl Fuerlinger, Ludwig Maximilian University Munich

W5-1 *Invited Keynote Talk*

Building Specialized Tools Using Tool Component Frameworks

Martin Schulz, Lawrence Livermore National Laboratory

W5-2 Critical-path-guided Interactive Parallelisation

Jonathan Mak and Alan Mycroft

W5-3 Pre-computing Function Results in Multi-Core and Many-Core Processors

Edward C. Herrmann, Prudhvi Jang and Philip A. Wilsey

Session II

Time: 15:20-16:20, September 13 (Tuesday)

Room: B2-3

Session Chair: Karl Fuerlinger, Ludwig Maximilian University Munich

W5-4 Assessing the Performance of MPI Applications Through

Time-Independent Trace Replay

Frederic Desprez, George S. Markomanolis, Martin Quinsony and Frederic Suter

W5-5 Simulation of Large-Scale HPC Architectures

Ian S. Jones and Christian Engelmann

W5-6 Scalable Control and Monitoring of Supercomputer Applications using an Integrated Tool Framework

Gregory R. Watson, Wolfgang Frings, Claudia Knobloch, Carsten Karbach and Albert L. Rossi

W6: Applications of Wireless Ad Hoc and Sensor Networks (AWASN)

Session I: WiMAX

Time: 09:00-10:00, September 13 (Tuesday)

Room: B2-4

Session Chair: Yuh-Shyan Chen

W6-1 Handling the Backhaul Link Failure Problem for Femto ABSs in IEEE 802.16m Environments

Yu-Chan Lin, Whai-En Chen and Meng-Hsuan Lin

W6-2 A Study for Connection Establishment in Femtocell Network

Yu-Ching Hsu, Show-Shiow Tzeng and Ching-Wen Huang

W6-3 Channel-aware Slot Assignment in OFDMA-based Mobile WiMAX Networks

I-Shyan Hwang, Bor-Jiunn Hwang and Chien-Yao Chiu

W6-4 A Femtocell-Assisted Data Forwarding Protocol in Relay Enhanced LTE Networks

Yuh-Shyan Chen, Chao-Chun Li and Wen-Lin Chiang

Session II: Wireless Sensor Networks

Time: 10:20-12:00, September 13 (Tuesday)

Room: B2-4

Session Chair: Tomotaka Wada

W6-5 A Controller-Assisted Distributed (CAD) Load Balancing Scheme for ZigBee Networks

Kuei-Li Huang, Chien-Chao Tseng, Jui-Tang Wang and Tsung-Hsi Yang

W6-6 Time-Synchronized versus Self-Organized K-Coverage Configuration in WSNs

Meng-Chun Wueng, Prasan Kumar Sahoo and I-Shyan Hwang

W6-7 i-Mace: Protecting Females from Sexual and Violent Offenders in a Community via Smartphones

Jou-Chih Chang, Pi-Shih Wang, Kang-Hsuan Fan, Shih-Rong Yang, De-Yuan Su, Min-Shiung Lin, Min-Te Sun and Yu-Chee Tseng

W6-8 A Lightweight Secure Data Aggregation Protocol for Wireless Sensor Networks

Hung-Min Sun, Chiung-Hsun Chen and Po-Chi Li

W6-9 Barrier Coverage Constructions for Border Security Systems using Wireless Sensors

Koji Yamamoto, Hayato Ozaki, Takuya Suzuki, Tomotaka Wada, Koichi

Mutsuura and Hiromi Okada

Session III: Mobile Networks and Services

Time: 13:30-14:50, September 13 (Tuesday)

Room: B2-4

Session Chair: Chih-Wei Yi

W6-10 End to End Security and Path Security in Network Mobility

Long-Sheng Li, Shr-Shiuan Tzeng and Rui-Chung Bai

W6-11 Connectivity Modeling of Vehicular Ad Hoc Networks in Signalized City Roads

Prasan Kumar Sahoo, Ming-Jer Chiang and Shih-Lin Wu

W6-12 From Spatial Reuse to Transmission Power Control for CSMA/CA Based Wireless Ad Hoc Networks

Han-Chiuan Luo, Eric Hsiao-Kuang Wu and Gen-Huey Chen

W6-13 A Runtime Partitioning Technique for Mobile Web Services

Muhammad Asif and Shikaresh Majumdar

W6-14 Routing and Buffering Strategies in Delay-Tolerant Networks: Survey and Evaluation

Shou-Chih Lo, Min-Hua Chiang, Jhan-Hua Liou and Jhih-Siao Gao

Session IV: Localization and Detection

Time: 15:20-17:00, September 13 (Tuesday)

Room: B2-4

Session Chair: Jehn-Ruey Jiang

W6-15 AR-based Positioning for Mobile Devices

Yaun-Chou Cheng, Ju-Yi Lin, Chih-Wei Yi, Yu-Chee Tseng, Lun-Chia Kuo, Yu-Jung Yeh and Chung-Wei Lin

W6-16 Parallel Response Query Tree Splitting for RFID Tag Anti-Collision

Ming-Kuei Yeh, Jehn-Ruey Jiang and Shing-Tsaan Huang

W6-17 An Accurate GPS-based Localization In Wireless Sensor Networks: A GM-WLS Method

Bo Cheng, Rong Du, Bo Yang, Wenbin Yu, Cailian Chen and Xinping Guan

W6-18 Sliding-Typed Communication Range Recognition Method for Indoor Position Estimation in Passive RFID Systems

Atsuki Inada, Yuki Oda, Emi Nakamori, Manato Fujimoto, Tomotaka Wada, Kouichi Mutsuura and Hiromi Okada

W6-19 Optimal Multipath Planning for Neyman-Pearson Detection in Wireless Sensor Networks

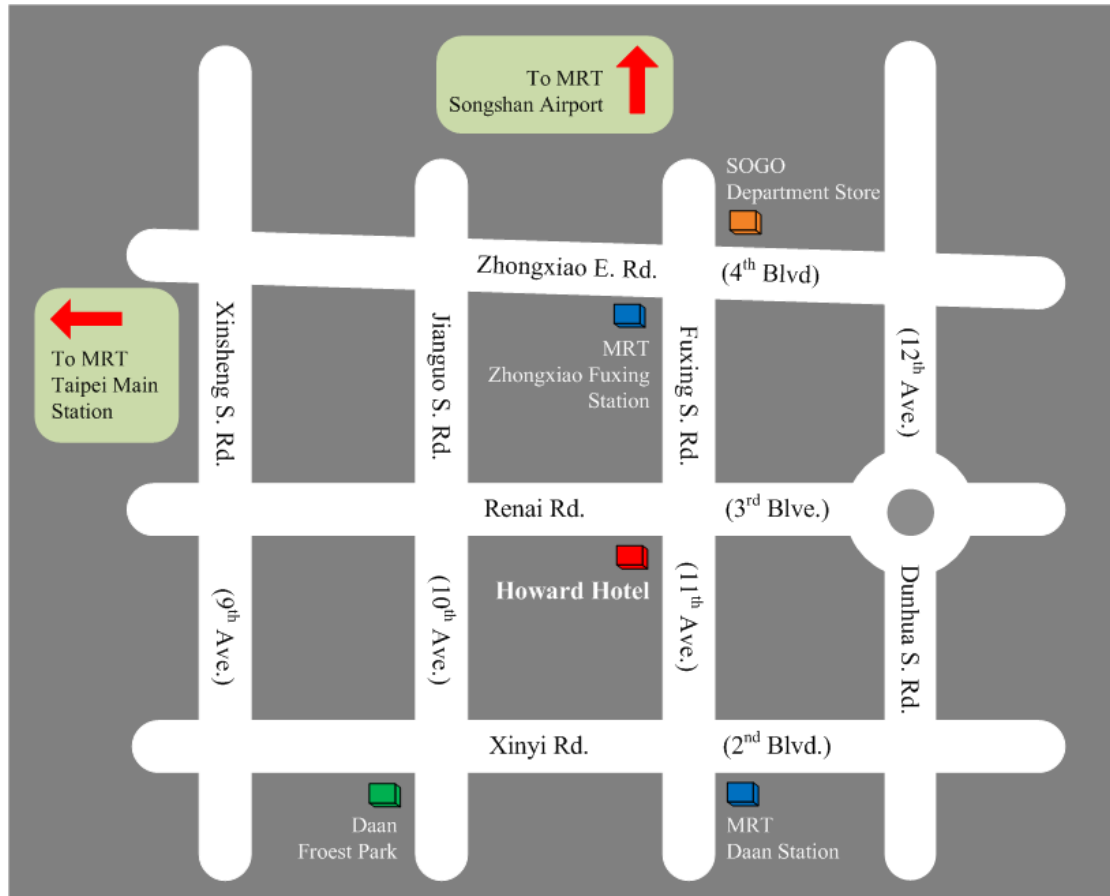
Yung-Liang Lai and Jehn-Ruey Jiang

Howard Hotel

Both the main conference and the workshop will be held on Basement 2.

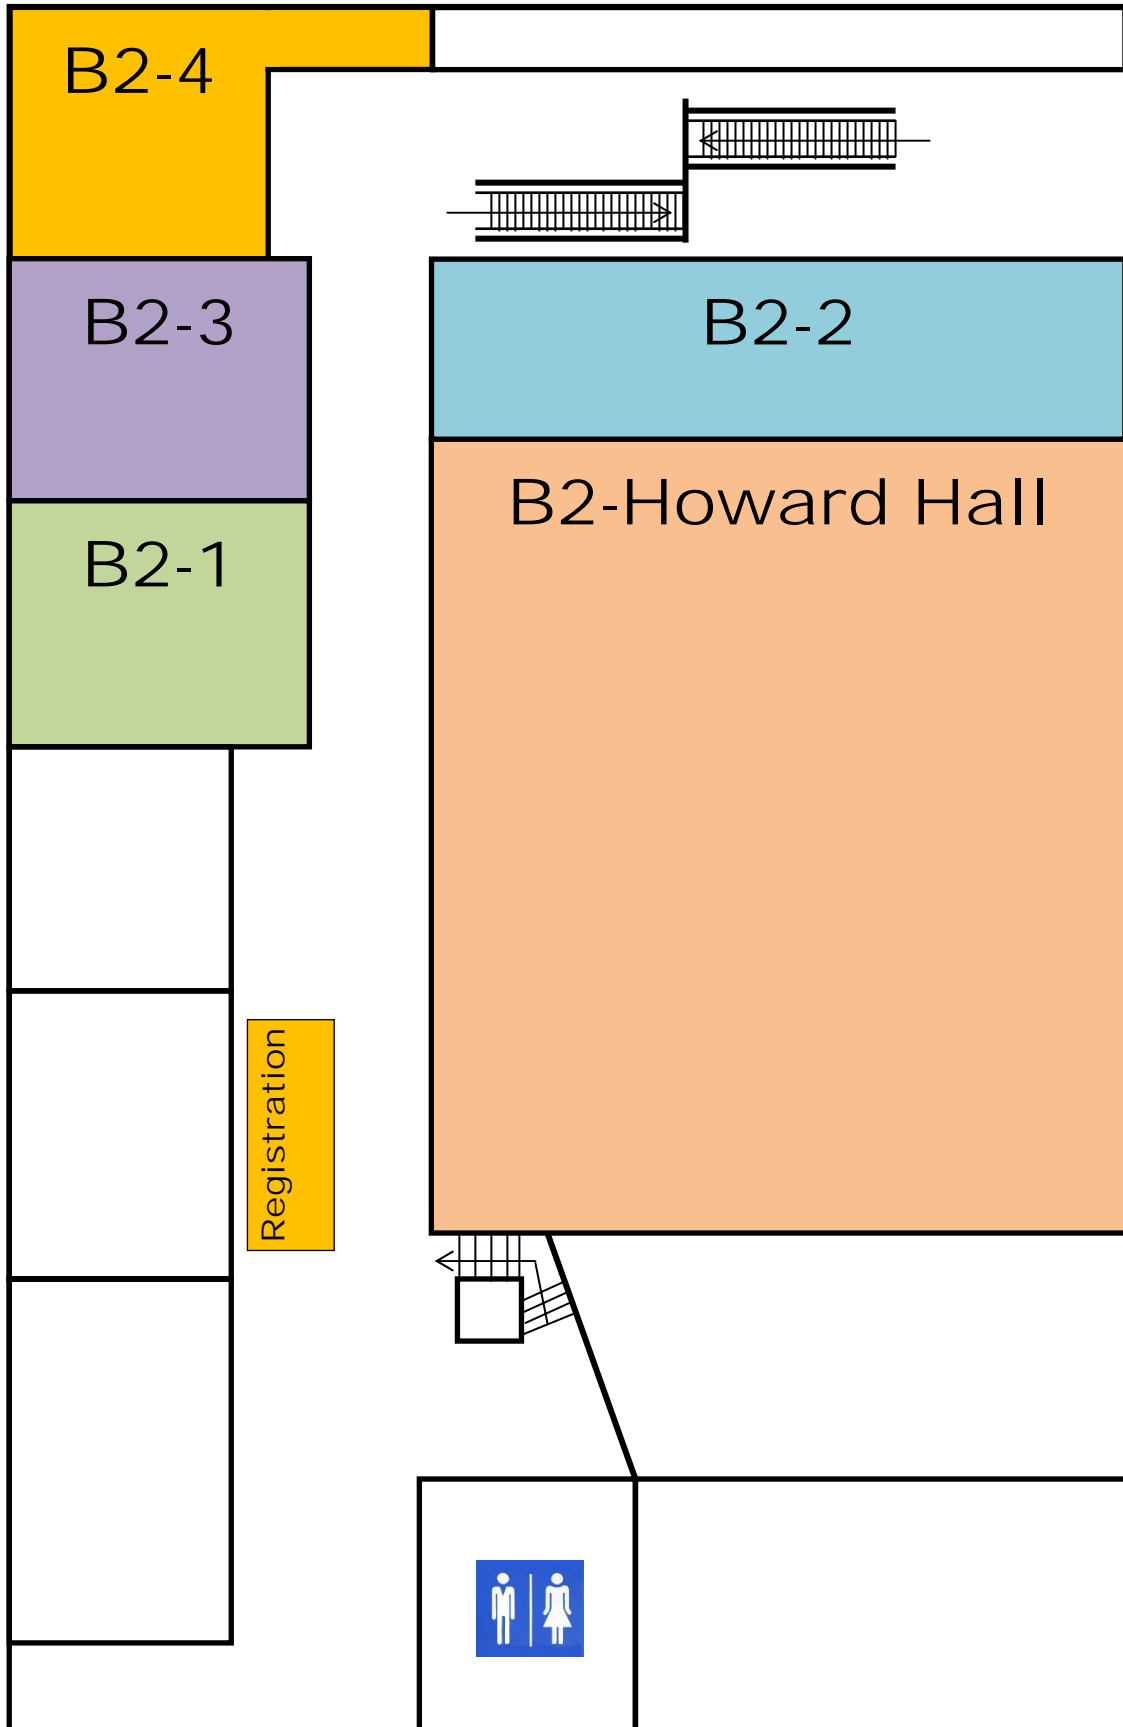
Address: No.160, Ren Ai Road Sec.3, Taipei, Taiwan

Telephone: 886-2-2700-2323

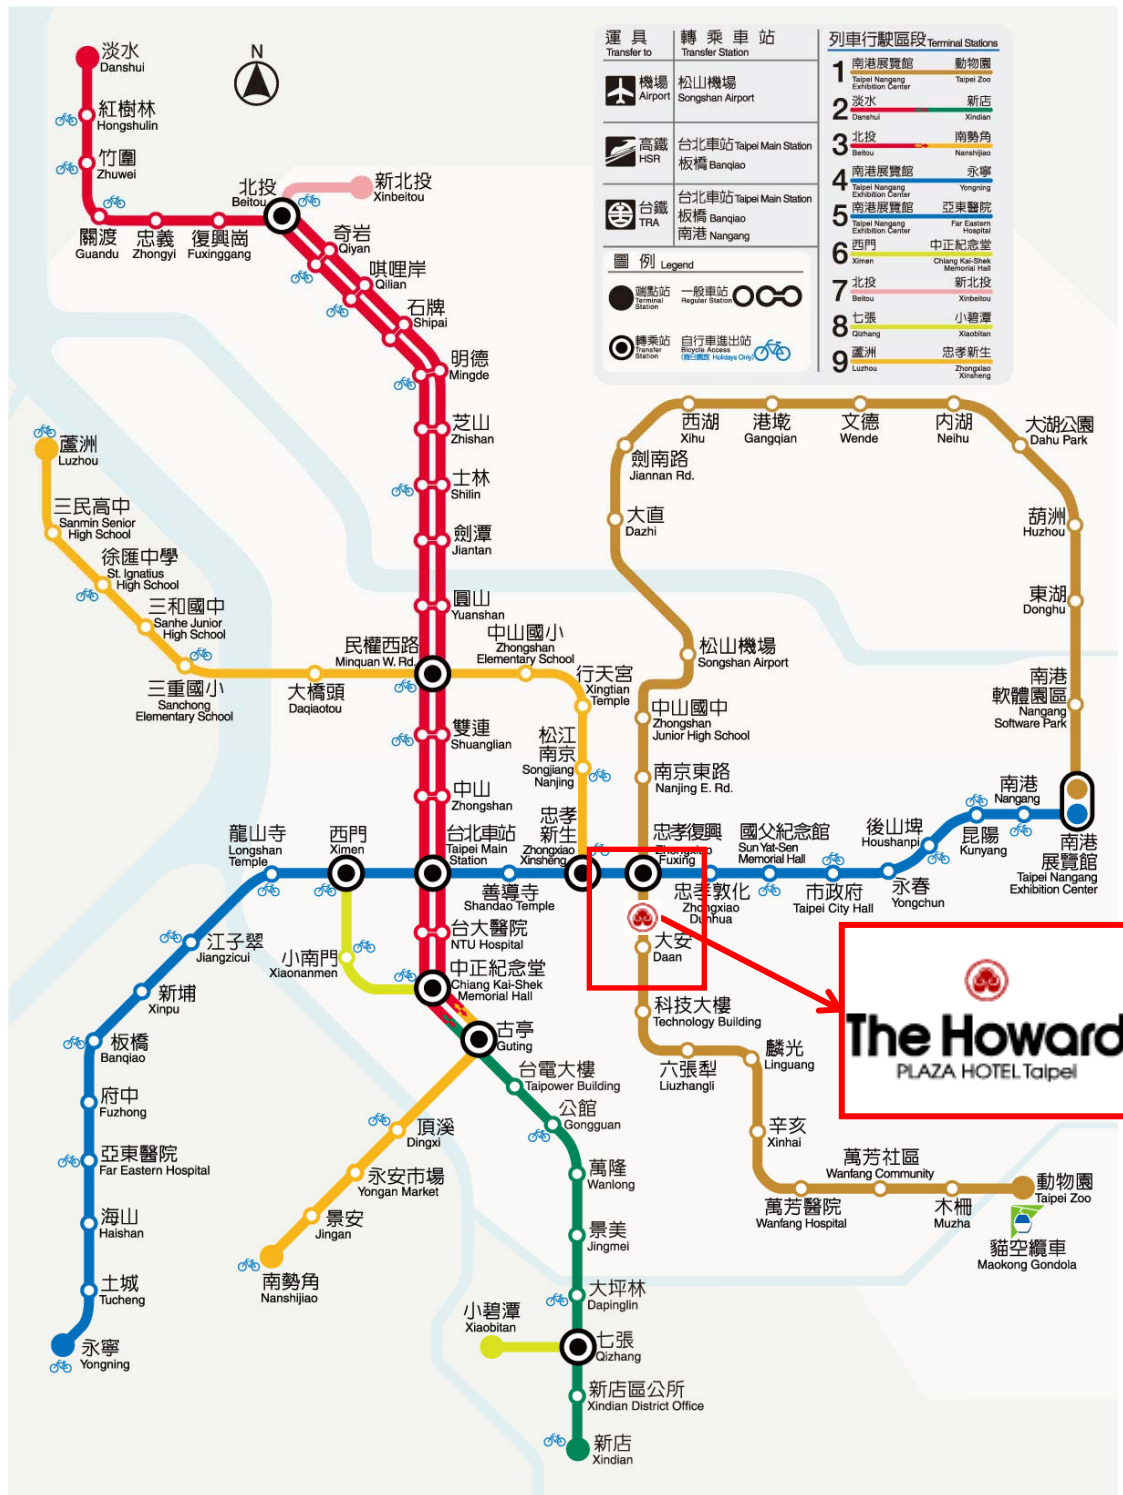


Locating in the prosperous east of Taipei city, the Howard Hotel Taipei provides you with luxurious accommodation combined with all the comforts of home in the capital city, and the decorating pays homage to classic oriental aesthetics. It is just steps to the shopping centers, chic restaurants, and all the attractions of east Taipei.

Map of Basement 2, Howard Hotel



Taipei MRT Route Map



Welcome Reception

Date: September 14 (Wednesday)

Time: 18:00-20:00

Restaurant: Rainbow Terrace the Atrium (彩虹座)


Location: 1F of Howard Hotel

Conference Banquet

Date: September 15 (Thursday)

Time: 18:00-20:30

Location: B2-Howard Hall, Howard Hotel



Conference Lunch

As part of the conference package, lunch has been arranged for full registered participants at three restaurants located on 2F and 3F of Howard Hotel: Champs Elysees (麗香苑), Pearl River (珍珠坊), and Yangtse River (江南春). You may use your lunch tickets, choose the restaurant of your preference and enjoy dining in any of the about three restaurants.

Date: September 13-16

Time: 12:00-13:30

Restaurant		Description
Champs Elysees (麗香苑)		2F of Howard Hotel Savor the rich array of eastern and western gourmet creations in a stylish environment with impeccable service.
Pearl River (珍珠坊)		2F of Howard Hotel Exquisitely prepared authentic Cantonese specialties and dim sums are perfectly complemented by superlative service for an unforgettable dining experience.
Yangtse River (江南春)		3F of Howard Hotel Well celebrated for its traditional Shanghainese cuisine with authentic flavor, Yangtse River promises a memorable dining experience.

Hosts

- National Chiao Tung University
- National Tsing Hua University
- Chung Yuan Christian University

Sponsors

- Bureau of Foreign Trade
- Chunghwa Telecom
- Department of Information and Tourism, Taipei City Government
- Industrial Technique Research Institute (ITRI)
- Institute for Information Industry (III)
- Institute of Information Science, Academia Sinica (IIS)
- Intel Corporation
- Microsoft Research Asia
- Ministry of Education
- National Science Council
- NCTU/ITRI Joint Research Center
- Research Center for Information Technology Innovation, Academia Sinica (CITI)
- Taiwan Association of Cloud Computing

