2017 WORKSHOP ON Advanced Computing and Networking

December 15-16, 2017

Room A101,School of Data & Computer Science

SYSU,Guangzhou.

HOST:

School of Data and Compter Science, SYSU

PROGRAM 会议日程表

时 间：2017 年 12 月 15 日-16 日（签到时间：12 月 15 日 8:30-9:00）

地 点：中山大学东校区 数据科学与计算机学院 A101

主 持：沈鸿教授

2017 年 12 月 15 日		
时间	报告题目	演讲者/单位
9:00-10:00	Scheduling of AI Tasks at the Edge	Professor Francis Lau, The University of Hong Kong
10:00-10:30	Tea Break 茶歇	
10:30-11:30	Low Resolution Face Image Recognition Using Deep Convolution Neural Networks	Professor ShiJinn Horng, National Taiwan University of Sci & Tech, Taiwan
11:30-12:30	Privacy-Preserving Computation on Network Data	Dr. Hui Tian, Univ. of Adelaide / BJTU
12:30-13:30	Lunch 午餐	
14:00-15:00	面向大规模多模态机器学习的异构并行处理方法、平台与应用	李肯立 教授 湖南大学
15:00-15:30	Tea Break 茶歇	
15:30-16:30	Optimizing MapReduce Framework in Multi-GPU systems	Professor Kuan-Ching Li, Providence University,Taiwan
16:30-17:30	Efficient Approximation Algorithms for Multi-Antennae Largest Weight Data Retrieval	Dr. Longkun Guo, Fuzhou University

PROGRAM 会议日程表

时 间：2017 年 12 月 15 日-16 日（签到时间：12 月 15 日 8:30-9:00）

地 点：中山大学东校区 数据科学与计算机学院 A101

主 持：沈鸿教授

2017 年 12 月 16 日		
时间	报告题目	演讲者/单位
9:00-10:00	工业物联网资源访问与节能技术	马礼 教授 北方工业大学
10:00-10:30	Tea Break 茶歇	
10:30-11:30	大数据泛构：应对多样性挑战	毛睿 教授 深圳大学
11:30-12:00	Efficient Design of Optical Network-on-Chips from a Networking Perspective	Dr. Yawen Chen, Otago University
12:30-13:30	Lunch 午餐	
14:00-15:00	脑科学中的高性能计算问题	冯圣中 研究员 中国科学院深圳先进技术研究院
15:00-15:30	Tea Break 茶歇	
15:30-16:30	基于超算平台的深度学习技术	吴维刚 教授 中山大学
16:30-17:30	Cyber-physical smart grid security	Dr. Neetesh Saxena, Bournemouth University

Scheduling of AI Tasks at the Edge

Professor Francis Lau, The University of Hong Kong

Abstract:

AI is now perceived to be a panacea for many real-life problems that can be solved computationally. It surely could have that ability, but most AI programs however operate in a mode that is rather resource-intensive and thus tend to be sluggish if run on ordinary computing devices. If AI has to work effectively in the front line, such as in an auto-driving vehicle, a checkpoint where the human face is used as ID, the cockpit of an airplane, or a power plant where unusual surges need to promptly attend to, it needs to run efficiently and in an interactive fashion. Traditionally, an AI program serving some device either runs on the device itself, using a much stripped down model of a neural network, which suffers from low accuracy results, or is relegated to a remote cloud, which takes an intolerably long time to respond because of the network distance. The situation had called for a new additional layer that is situated in close proximity to the devices (the "things" in the much trumpeted Internet of Things), which is called the ***edge***. Now AI tasks can choose to run in the edge where there are more powerful servers which are at just an arm's length away; the result is better accuracy at interactive or close to interactive speed. But who or how to make that decision to choose the edge (and why not the thing itself or the cloud?) is a critical question when operating within this new trichotomy in real life. The decision has to come from some intelligent software: an algorithm that can, given an AI task that can be solved using one of many possible models (corresponding to different accuracies and resource requirements) and its time constraint, decide the best choice for the task as well as for the system in terms of effective usage of resources. The design of such an algorithm becomes a challenge when it has to handle multiple tasks originating from different users at the same time. This talk will discuss possible approaches to meeting this challenge and outline the corresponding designs of workable algorithms for handling the said tasks.

Francis Lau is a professor in computer science and associate dean of Engineering at The University of Hong Kong. He served as the head of the Department of Computer Science from 2000-2005. From 1993 to 2001, he was actively involved in many activities under IEEE Computer Society, which culminated at him serving as a vice president (for chapters activities) of the society in 1999. He was part of the team that drafted the IEEE/ACM Computer Curricula 2001 which had influenced the design of many CS curricula around the world. Professor Lau's research interests include parallel & distributed systems, algorithms, programming languages, AI, and application of computing to music and art, etc. He is the editor-in-chief of the Journal of Interconnection Networks (JOIN) published by World Scientific. In his spare time, Professor Lau indulges in classical music.

Low Resolution Face Image Recognition Using Deep Convolution Neural Networks

洪西进教授，台湾科技大学

Abstract:

Deep leaning is a very useful tool for many applications especially for pattern recognitions and computer vision. Low resolution face recognition is one of challenging problem in computer vision. Due to the difficulty of finding the specific features of faces, the accuracy of recognition is still quite low. In this talk we will introduce how to solve this problem using Deep learning techniques. The proposed method consists of two major parts; one is Restricted Boltzmann Machine and the other is Deep Convolution Neural Network. The former is used to preprocess the face images and the latter is used to do classification. Based on the existing databases, we then conduct the training and testing. Compared to existing methods, the proposed method can really improve the accuracy of recognition. Then the face recognition technique can be further applied widely and efficiently.

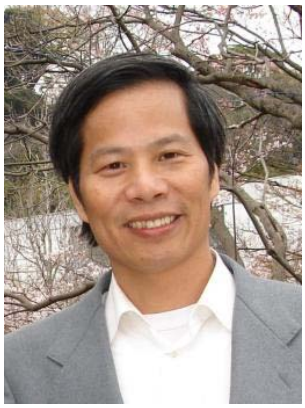

洪西进，教授，男，工学博士。目前是台湾科技大学的讲座教授。洪教授曾担任台湾科技大学资讯工程系主任(2013~2016)，计算器中心主任和联合大学电资学院院长(2006~2009)。2012 年四川省“百人计划”入选者，2013 年四川省科技顾问。分别于 2010 年，2008 年，2004 年担任美国 Georgia State University，日本 Tokyo Institute of Technology，蒙古 University of Mongolia 的访问教授。 洪教授曾担任 9 个国际会议的会议主席，亦曾受邀担任 5 个国际会议的 Keynote Speaker。洪教授已经发表了超过 200 篇研究论文，并获得许多奖项，如 2011 年 *IEEE Transactions on Industrial Informatics* 期刊的最佳论文奖， 2006 年中国电机工程学会“杰出电机工程教授奖”，2004-2007 年国科会“杰出研究奖”等。洪教授的研究兴趣包括生物辨识、信息安全、超大规模集成电路设计、多处理器系统和并行算法。

面向大规模多模态机器学习的异构并行处理方法、平台与应用

李肯立教授，湖南大学

摘要：

多模态机器学习是当前人工智能领域的研究热点，具有数据量大、过程复杂、模态之间需多层次协同耦合等特点，对高性能计算提出了巨大的需求和挑战。将深入分析大规模多模态机器学习并行处理方法及平台（如 Tensorflow 等）的国内外研究现状与不足，结合我国国产高性能计算机体系结构方面研究的最新进展，遵循“大规模多模态机器学习处理分析与建模-并行处理方法与平台 - 核心可扩展并行算法设计-算法及平台稳定性与优化-在铁路货运列车故障检测中的应用”的研究路线，从多模态数据机器学习建模的 核心算法和操作出发，论述基于国产自主异构众核系统的机器学习并行处理新方法研究的必要性，可能存在的理论和技术难点和挑战，及其可能的应用领域。

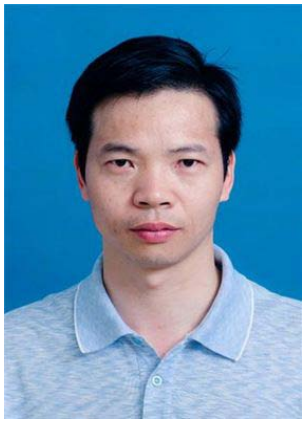

李肯立，现任湖南大学信息科学与工程学院教授、博士生导师、院长，国家超级计算长沙中心主任，数据分析湖南省工程研究中心主任。2003 年博士毕业于华中科技大学国家高性能计算中心（武汉）。先后获得国家杰出青年基金、科技部中青年创新计划领军人才，湖南省芙蓉学者特聘教授等人才计划支持，IEEE Senior Member，CCF 杰出会员，超级计算创新联盟副理事长。 主要研究领域为并行与分布式处理、高性能计算和大数据管理等，担任 IEEE-TC、IJPRAI 等杂志编委，先后在 IEEE-TC、IEEE-TPDS、IEEE-TSP、IEEE-TIFS 和 ICPP、ICDCS 等国际顶级期刊和会议发表学术论文 170 余篇，其中 IEEE/ACM Transaction 期刊近 50 篇。主持国家自然科学基金重点项目、科技部国际科技合作专项项目等国家和省部级项目 20 余项，获湖南省科技进步一等奖和教育部科技进步二等奖等奖励 4 项。

脑科学中的高性能计算问题

冯圣中研究员，中国科学院深圳先进技术研究院

摘要：

高性能计算技术和脑结构、脑信号实验观测手段的迅速发展，为逐步揭开蒙在大脑上的神秘面纱提供了可能。构成脑的基本结构单元是数亿甚至数千亿的神经元，神经元突触连接构成超大规模网络。分析精细脑结构-功能关系，及大脑网络的结构和连接特征，对于揭示脑功能机理具有重要理论意义与应用价值。这些分析工作也对高性能计算提出系列挑战。本报告将在脑网络特点、复杂网络分析以及复杂脑网络研究现状三方面介绍的基础上，分析脑科学中的高性能计算问题。

冯圣中，中国科学院“百人计划”研究员，博士生导师，中国计算机学会杰出会员，科技部高性能计算重点领域专家，中国科学院深圳先进技术研究院院长助理。主持或参与国家863、国自然重点、中科院知识创新工程、广东省重大技术攻关、深圳市重大技术攻关等项目。荣获中国科学院杰出贡献奖、国家科技进步二等奖、“春晖杯”海外留学人员创新创业大赛一等奖、微软 HPC 学生教师研究奖第一名等奖项。长期从事高性能计算、大数据技术、生物信息学等方向研究及产业化方面工作。

Optimizing MapReduce Framework in Multi-GPU Systems

Professor Kuan-Ching Li, Providence University, Taiwan

Abstract:

The potential of Big Data is widely recognized and organizations are collecting huge amounts of data. Most of them have the same goal: to extract "value" through sophisticated analysis for findings and decisions. Rapid advancements in computer and networking technologies have been critical to the IT revolution, driving the rapid decline in the cost and fast increase in the processing power of digital technologies. Concurrently, in the ever-changing world of software development, it's extremely important to keep up with current technologies, methodologies, and trends. MapReduce is a programming model introduced by Google for large-scale data processing, and a number of studies have implemented MapReduce model on GPUs. However, most of them are based on the single GPU and bounded by GPU memory with inefficient atomic operations. In this presentation, we present a standalone MapReduce system to utilize multiple GPUs, handle large-scale data processing beyond GPU memory limit, and eliminate serial atomic operations. Experimental results have demonstrated the proposed research's effectiveness and promising the handling of large datasets.

Kuan-Ching Li is a Distinguished Professor of computer science and engineering at Providence University, Taiwan. He is a recipient of guest and distinguished chair professorships from universities in China and other countries, and awards and funding support from a number of agencies and industrial companies. He has been actively involved in many major conferences and workshops in program/general/steering conference chairman positions, and has organized numerous conferences related to high-performance computing and computational science and engineering. He is a Fellow of IET, senior member of the IEEE and a member of the AAAS, editor-in-Chief of International Journal of Computational Science and Engineering (IJCSE), International Journal of Embedded Systems (IJES), and International Journal of High Performance Computing and Networking (IJHPCN), published by Inderscience. Besides publication of research papers, he is co-author/co-editor of several technical professional books published by CRC Press, Springer, McGraw-Hill and IGI Global. His research interests include GPU/many-core computing, Big Data, and cloud.

工业物联网资源访问与节能技术

马礼教授，北方工业大学

摘要:

工业物联网迅速发展，其安全可靠接入和访问至关重要，利用节点 IP 地址进行管理仍然是主要手段。IP 地址具有标志节点身份和位置信息的两项基本功能，特别是无线移动节点，其拓扑和位置经常改变，真实表示和记录节点身份与位置，至关重要。从工业物联网体系结构出发，研究资源对象命名服务、域名分配与寻址、分布式映射表项存储与管理、安全访问等关键问题。通过实践，解决工业物联网全网节点实名登记和认证，运行安全可靠等问题。研究中利用 IPv6 作为节点地址分配基本结构，综合目前工业物联网已经形成的三大标准，即 WirelessHART、ISA100.11a、WIA-PA，设计异构接口一致访问协议，通过嵌入语义的地址编码与分配、统一地址映射与访问等问题，有效解决实名可靠接入、跟踪节点主要访问行为、确保访问可靠等关键问题。

马礼，教授，博士。北京市教学名师，北京市高校学术创新团队带头人，中国计算机学会理事，北京市高等教育学会计算机教育研究会理事，CCF 物联网专委、教育专委会委员，IEEE-CS 会员、ACM 会员。研究方向包括分布式信息处理、高性能计算、多 agent 机器人系统等，先后主持完成了国家科技支撑计划、北京市自然科学基金等项目 20 多项，在分布式系统领域有丰富的研究经历。完成了大气与环境科学重点学科网格系统建设、基于 MAS 的多机器人系统合作与协调技术等研究，实现了无线传感器网络自主组网，异构多机器人合作与协调系统等原型机项目。获得全军科技进步奖励 1 项，获得省级教学成果二等奖 2 项，公开发表学术论文 80 多篇。主编并出版了《计算机组成原理与系统设计》等教材 4 部。

大数据泛构：应对多样性挑战

毛睿教授，深圳大学

摘要:

目前的大数据处理主要都围绕大数据的 volume 和 velocity 这两个挑战，对于 variety 挑战的研究相对较少。通用的数据处理技术因其广泛的适用性和相对低的开发成本，一直受到商业数据库管理系统的亲睐。大数据泛构（big data genhierarchy）把数据抽象成度量空间中的点，利用数据间距离的三角不等性进行数据的管理和挖掘等工作，具有高度的通用性。基于这一思想，我们开发了 UMAD 软件包，构建通用大数据管理挖掘框架。

毛睿，男，博士，教授，主要研究方向通用大数据管理分析方法和高性能计算；97 年和 00 年在中国科学技术大学计算机系获学士和硕士学位，06 年和 07 年在美国得克萨斯大学奥斯汀分校获统计学硕士和计算机科学博士学位，07～10 年在甲骨文美国公司任高级工程师；10 年加入深圳大学，现任计算机与软件学院副院长、大数据系统计算技术国家工程实验室副主任；CCF 高级会员、数据库专委常委、大数据专家委委员、高性能计算和理论计算机科学专委委员。

基于超算平台的深度学习技术

吴维刚教授，中山大学

摘要：

随着大数据时代的到来，各类大数据应用进入快速发展的时期。纷繁多样的数据和千差万别的应用需求给大数据处理带来了很大挑战。基于简单集群的大数据分析处理平台和技术已经难以满足智能化、深层次的数据分析处理需求。特别是深度学习技术，由于深度神经网络的规模越来越大，对计算资源的需求也快速增长。而超级计算系统有充沛的计算资源，可以为大数据处理提供有力支撑。本报告聚焦基于超级计算平台的深度学习算法的并行优化问题，讨论其中的关键技术难点，并介绍天河二号系统上开展的深度学习并行优化方面的初步探索。

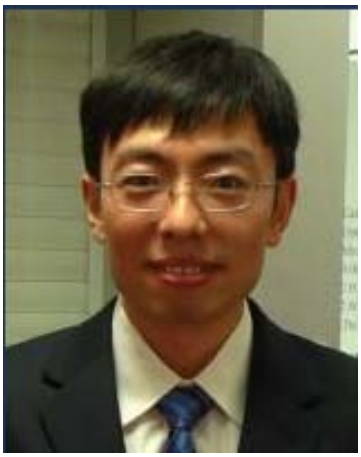

吴维刚，教授、博导，中山大学数据科学与计算机学院院长助理、超算专项工作办公室主任，广州市超算与大数据重点实验室副主任。分别于 1998 年和 2003 年在西安交通大学获得学士和硕士学位；于 2007 年获香港理工大学计算机博士学位。于 2008 年通过中山大学“百人计划”加入中山大学计算机系。2011 年获广州市“珠江科技新星”计划资助。主要从事分布式计算与网络方面的研究工作，研究方向包括：云计算资源管理、大数据管理与处理、车联网与边缘计算等。已在国内外发表学术论文 70 余篇，包括 IEEE TPDS、IEEE TC、《软件学报》、《中国科学》、IEEE SRDS、IEEE ICPP 等权威期刊和会议论文十多篇。获得 2015 年广东省计算机学会优秀论文一等奖。主持国家重点研发计划课题、国家自然科学基金项目、广东省自然科学基金项目等各类科研项目十余项。目前担任两个 SCI 期刊 Frontiers of Computer Science 和 Ad Hoc & Sensor Wireless Networks 的编委。担任 CCF 分布式计算与系统、系统软件、普适计算专委会委员、CCF 广州分部委员。

Privacy-Preserving Computation on Network Data.

Dr. Hui Tian, Univ. of Adelaide / Beijing Jiaotong Univ.

Abstract:

In the era of cloud computing with the evolving demand of big data processing, privacy-preserving computing (PPC) has arisen as an effective way to achieve secure distributed computing and information sharing which serves as the base for realization of widespread Smart City and e-Society. PPC requires to develop a computation paradigm for solving a given problem that takes privacy-protected data as input and produces an output that is utilizable to the public yet secure against privacy attacks. For example, in a health network, patients' data from different hospitals and health agents may be analyzed for discovery of disease-district and disease-habit relationships. How to protect the privacy of patients while maintaining a certain degree of data utility is a typical PPC problem for network traffic data. In social networks, the number of interacting parties represented by node degree, the frequency of two interacting nodes described by their link's weight are also private information. How to protect these information while publishing utilizable social network structures is a PPC problem for network topology data. In this talk, I will begin with an overview of existing techniques of privacy protection for network data publication. I will introduce our work on how to perform privacy preserving computation of set operations for big data analysis that enables us to find out intersection and matching tuples of multiple many-attribute data sets without revealing individual content. Secondly, I will show our recent work on privacy protection for social network topology data publishing by applying two different techniques of combining differential privacy with wavelet transform and k-anonymity respectively. Finally I will conclude the talk by presenting some interesting open problems.

Hui Tian is currently a Research Associate in University of Adelaide, and also Associate Professor in School of Electronics and Information Engineering, Beijing Jiaotong University (BJTU). She received B. Eng. and M. Eng. degrees from Xidian University, China and Ph.D. from Japan Advanced Institute of Science and Technology. She was an Associate Professor (Research Fellow) in Institute of Computing Technology of Chinese Academy of Science during 2007-2009, and a Lecturer in Manchester Metropolitan University, UK, during 2005-2006. She has published over 40 papers in international journals and conferences. Her research interests include network performance evaluation, telecommunications, privacy preserving computing and wireless sensor networks.

Efficient Approximation Algorithms for Multi-Antennae Largest Weight Data Retrieval

Dr. Longkun Guo, Fuzhou University

Abstract:

In a mobile network, wireless data broadcast over channels (frequencies) is a powerful means for distributed dissemination of data to clients who access the channels through multi-antennae equipped on their mobile devices. The δ -antennae largest weight data retrieval (δ ALWDR) problem is to compute a schedule for downloading a subset of data items that has a maximum total weight using δ antennae in a given time interval. In this paper, we first give a linear programming (LP) relaxation for δ ALWDR and show that it is polynomial-time solvable when every data item appears at most once. We also show that when there exist data items with multiple occurrences, the integrality gap of this LP formula is 2 . We then present an approximation algorithm of ratio 1- 1/e for the δ-antennae γ-separated largest weight data retrieval (δA-γ-LWDR) problem, a weaker version of δ ALWDR where each block of up to γ data (time) slots is separated by a vacant slot on all channels, applying the techniques called collectively randomized LP rounding and layered DAG construction. We show that δ A γ LWDR is NP-complete even for the simple case of γ =2 , m=3 , and equal-weight data items each appearing up to 3 times. Then, from the simple observation that a ratio α approximation solution to δA-γ-LWDR implies a ratio α-ε approximation solution to δ ALWDR for any fixedε>0, we immediately have an approximation algorithm of ratio 1- 1/e-εfor δ ALWDR.

Longkun Guo received the B.S. and ph.D. degrees in Computer Science from University of Science and Technology of China (USTC). He was a research associate of the University of Adelaide, and is currently an associate professor and the associate head of the department of computer science, College of Mathematics and Computer Science, Fuzhou University. His major research interest is efficient algorithm design and computational complexity analysis for optimization problems in high performance networks. Longkun has published more than 30 academic papers, many of which are in top-tier academic journals/conferences, including A*/A journals/conferences ranked by the CCF, ERA and/or CORE, such as Algorithmica, IEEE Transactions on Mobile computing, ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), etc. He served on the technical program committees of high performance computing and/or networking conferences including PDCAT, CloudCom-Asia, PAAP, etc.

Efficient Design of Optical Network-on-Chips from a Networking Perspective

Dr. Yawen Chen, Otago University

Abstract:

Nowadays microprocessor development has moved into a new era of many-core on-chip design, with tens or even hundreds of cores fitting within a single processor chip to speed up computing. However, conventional electrical interconnect for inter-core communication is limited by both bandwidth and power density, which creates a performance bottleneck for microchips in modern computer systems - from smartphones to supercomputers, and to large-scale data centers. Optical Network-on-Chip (ONoC), a silicon-based optical interconnection among cores at the chip level, overcomes the limitations of conventional electrical interconnects by supporting greater bandwidth with less energy consumption, and opens the door to bandwidth- and power-hungry applications such as big data analysis and general artificial intelligence. However, existing ONoC designs do not fully take the advantages of optical communication to maximize performance and save energy at the networking level. This talk will introduce our ongoing challenging research problems from a networking perspective and present our current results for designing efficient routing schemes specific for ONoCs. Other ongoing research topics in system research group of Otago University will also be briefly introduced.

Yawen Chen obtained her PhD degree in Computer Science from The University of Adelaide in Australia in 2008 under the supervision of Professor Hong Shen. After her Phd study, she worked as postdoctoral researcher at Royal Institute of Technology (KTH) in Sweden and Otago University in New Zeland during 2009-2011. She was appointed as Lecturer in Department of Computer Science at Otago University from 2012 and promoted to Senior Lecturer from 2018. She received Marsden Fund awarded by Royal Society of New Zealand (Principal Investigator, successful rate ~7%) in 2016. Her research interests include network-on-chips, optical networks, wireless optical networks, multicore computers, wired/wireless networks.

Cyber-Physical Smart Grid Security

Dr. Neetesh Saxena, Bournemouth University

Abstract:

It has been proved by the recent attacks on the power grid around the world that the smart grid is prone to the cyber attacks. In order to understand current health of the power system, we need a tool to simulate the attack environment and prepare the system for any real-time cyber attacks. In this talk, I will discuss a cyber-physical smart grid security assessment tool, explain its functional requirements, process execution, and overall working of the tool. At the end, I will also present a case study.

Neetesh Saxena is a Lecturer in Cyber Security with the Department of Computing and Informatics. Before joining BU, he was a Post-Doctoral Researcher with the School of Electrical and Computer Engineering at the Georgia Institute of Technology, USA. Prior to this, he was with the Department of Computer Science, The State University of New York (SUNY) Korea, South Korea as a Post-Doctoral Researcher and a Visiting Scholar at the Department of Computer Science, Stony Brook University. He earned his PhD in Computer Science and Engineering from the Indian Institute of Technology, Indore, India. In 2013-14, he was a Visiting Research Student and a DAAD Scholar at Bonn-Aachen International Center for Information Technology (B-IT), Rheinische-Friedrich-Wilhelms Universität, Bonn, Germany. He was also a TCS Research Scholar during Jan. 2012 - Apr. 2014. He works in the area of security and privacy. His current research interests include cyber security, cyber-physical system security in the smart grid and vehicle-to-grid, security and privacy in the cellular networks, securing end-to-end systems, and secure mobile applications. He is an editorial member of Springer Plus, and a reviewer of several international journals and conferences, such as IEEE TMC, IEEE TC, IEEE SJ, IEEE TII, IEEE TME, WPC Springer, FGCS Elsevier, IEEE WOCN, IEEE ICUFN, etc. He also serves as a TPC member of IEEE SmartGridCom'16, IEEE TrustCom'16, IEEE WiMob'16, IEEE PIMRC'16, IEEE/CIC ICCC'16, etc. He is a member of IEEE and ACM.