计算机科学最新进展高端研讨会

High-end Forum on Latest Development in Computer Science

June 19 2018
Conference Hall
Second Floor
National Supercomputing Center in Changsha

主办： “111计划” 高效能计算学科创新引智基地
湖南大学信息科学与工程学院
国家超级计算长沙中心

High-end Forum on Latest Development in Computer Science

Date: June 19, 2018, Tuesday, 8:30 am – 18:00 pm

Location: Conference Hall, National Supercomputing Center in Changsha, Hunan University, Changsha, Hunan, China.

Sponsors: "111 plan" High Efficiency Computing Innovation and Talent Base

College of Computer Science and Electronic Engineering at Hunan University

National Supercomputing Center in Changsha

Objectives: The purpose of this high-end forum is to bring distinguished scholars to share their latest development in all areas of computer science and to provide a unique opportunity for the faculty and students at Hunan University to communicate with leading researchers.

Forum Chair: Keqin Li (SUNY Distinguished Professor and Distinguished Professor of Chinese National Thousand Talents Program at Hunan University)

Program:

June 18, 2018 (Monday)

18:00-20:00 Welcome Banquet for Distinguished Speakers Host: **Kenli Li** (Dean of College of Information Science and Engineering at Hunan University; Director of National Supercomputing Center in Changsha)

June 19, 2018 (Tuesday)

8:30-8:45
Opening Remark: **Weihong Tan** (Vice President for International Affairs of Hunan University, Academician of Chinese Academy of Science)

8:45-9:30
Distinguished Speech 1: On System Challenges for Smart Machines, **Guang R. Gao** (Distinguished Professor at the University of Delaware, ACM Fellow and IEEE Fellow)

9:30-10:15
Distinguished Speech 2: An Introduction to AI and Deep Learning, **John E. Hopcroft** (IBM Professor of Engineering and Applied Mathematics in Computer Science at Cornell University, A. M. Turing Award Winner, Member of the US National Academy of Sciences and the US National Academy of Engineering)

10:15-10:30
Coffee Break

10:30-11:15
Distinguished Speech 3: Reconstructing Reality: From Physical World to Virtual Environments, **Ming C. Lin** (Elizabeth Stevinson Iribe Chair of Computer Science at the University of Maryland - College Park, ACM Fellow and IEEE Fellow)

11:15-12:00
Distinguished Speech 4: Accelerated Computing Using FPGAs, **Viktor K. Prasanna** (Charles Lee Powell Chair in Engineering at the University of Southern California, ACM Fellow and IEEE Fellow)

12:00-1:30
Lunch Break

1:30-14:15
Distinguished Speech 5: Supercomputing: Past, Present and Future, **Marc Snir** (Michael Faiman Professor at the University of Illinois at Urbana-Champaign, ACM Fellow and IEEE Fellow)

14:15-15:00
Distinguished Speech 6: Decomposed Fuzzy Systems, **Shun-Feng Su** (Chair Professor of Electrical Engineering at the National Taiwan University of Science and Technology, IEEE Fellow)

15:00-15:15
Coffee Break

15:15-16:00
Distinguished Speech 7: Using Kernels to Harness the Complexity of Big Data Applications， **Benjamin W. Wah** (Provost and Wei Lun Professor at the Chinese University of Hong Kong, ACM Fellow and IEEE Fellow)

16:00-16:45
Distinguished Speech 9: Some Considerations of Network Computing Models, **Yaoxue Zhang** (Former President of Central South University, Academician of Chinese National Academy of Engineering)

16:45-17:30
Closing Remark: **Kenli Li** (Dean of College of Information Science and Engineering at Hunan University; Director of National Supercomputing Center in Changsha)

17:30-19:30
Banquet for Distinguished Speakers Host: **Kenli Li** (Dean of College of Information Science and Engineering at Hunan University; Director of National Supercomputing Center in Changsha)

On System Challenges for Smart Machines

Guang R. Gao
University of Delaware,
Distinguished Professor at the University of Delaware, ACM Fellow and IEEE Fellow

Abstract
We are facing the challenges of the end of Moore' s Law, as well as the challenges from applications in Smart Machines – propelled by advances in intelligent data analytics, machine learning and brain-inspired computing models. The speaker believes that it may be the time to initiate a new forum encouraging broader participation and direct interaction of scientists and engineers working on computer system architecture, system software (including compilers, runtime systems, and OS) and high level programming programing models and methodology, as well as scientists who are mapping demanding applications on to future smart machines. The talk will focused on the software foundation： program composability and scalability in large-scale parallel and distributed systems .
To this end, PTM – a parallel Turing Machine model is discussed as an attempt toward this direction (抛砖引玉！).

Biography
Guang R. Gao is a computer scientist and the 2017 Recipient of the IEEE Computer Society B. Ramakrishna Rau Award.
Gao has received his undergraduate education in Tsinghua University in Beijing. He received his Master and Ph.D degree in 1982 and 1986 respectively both in Computer Science at Massachusetts Institute of Technology (MIT) -- the first from mainland China. Gao has devoted his research and academic careers in dataflow model of computation To this end, Gao has led a series of parallel architecture and system projects where various aspects of dataflow models explored – ranging from innovations in programming paradigms, architecture features, and system software technology, including novel program optimization and runtime system techniques.
Gao is an ACM Fellow and IEEE Fellow. He has been awarded an Endowed Named Professorship of Electrical and Computer Engineering at University of Delaware. He received the Outstanding Achievement Award of Oversee Scholars - a prestigious award given by the CCF (Chinese Computer Federation). Gao is a founder and chairman of the IEEE/CS Dataflow STC – the Special Interested Community of Parallel Model and System: Dataflow and Beyond (http://dfstc.capsl.udel.edu/).

An Introduction to AI and Deep Learning

John Hopcroft
Cornell University
IBM Professor of Engineering and Applied Mathematics in Computer Science at Cornell University, A. M. Turing Award Winner, Member of the US National Academy of Sciences and the US National Academy of Engineering

Abstract

A major advance in AI occurred in 2012 when AlexNet won the ImageNat competition with a deep network. The success was sufficiently better than previous years that deep networks were applied in many applications with great success. However, there is little understanding of why deep learning works. This talk will give an introduction to machine learning and then illustratresearch e current directions in deep learning at a level for a general scientific audience.

Biography

John Hopcroft is the IBM Professor of Engineering and Applied Mathematics. He received his Ph.D. (1964) in EE from Stanford University. He chaired the Computer Science Department and was Dean of the College of Engineering. He is a member of the National Academy of Sciences, of the National Academy of Engineering, and a foreign member of the Chinese Academy of Sciences. In 1992, he was appointed by President George H. W. Bush to the National Science Board, which oversees the National Science Foundation, and served through May 1998. He has received numerous awards and honorary degrees. In 2016, Premier Li Keqiang presented him with the Friendship Award for his work in China.

Reconstructing Reality: From Physical World to Virtual Environments

Ming C. Lin
University of Maryland
Elizabeth Stevinson Iribe Chair of Computer Science at the University of Maryland - College Park, ACM Fellow and IEEE Fellow

Abstract

With increasing availability of data in various forms from images, audio, video, 3D models, motion capture, simulation results, to satellite imagery, representative samples of the various phenomena constituting the world around us bring new opportunities and research challenges. Such availability of data has led to recent advances in data-driven modeling. However, most of the existing example-based synthesis methods offer empirical models and data reconstruction that may not provide an insightful understanding of the underlying process or may be limited to a subset of observations.

In this talk, I present recent advances that integrate classical model-based methods and statistical learning techniques to tackle challenging problems that have not been previously addressed. These include flow reconstruction for traffic visualization, learning heterogeneous crowd behaviors from video, simultaneous estimation of deformation and elasticity parameters from images and video, and example-based multimodal display for VR systems. These approaches offer new insights for understanding complex collective behaviors, developing better models for complex dynamical systems from captured data, delivering more effective medical diagnosis and treatment, as well as cyber-manufacturing of customized apparel. I conclude by discussing some possible future directions and challenges.

Biography

Ming C. Lin is currently the Elizabeth Stevinson Iribe Chair of Computer Science at the University of Maryland College Park and John R. & Louise S. Parker Distinguished Professor Emerita of Computer Science at the University of North Carolina (UNC), Chapel Hill. She is also an honorary Chair Professor (Yangtze Scholar) at Tsinghua University in China. She obtained her B.S., M.S., and Ph.D. in Electrical Engineering and Computer Science from the University of California, Berkeley. She received several honors and awards, including the NSF Young Faculty Career Award in 1995, Honda Research Initiation Award in 1997, UNC/IBM Junior Faculty Development Award in 1999, UNC Hettleman Award for Scholarly Achievements in 2003, Beverly W. Long Distinguished Professorship 2007-2010, Carolina Women's Center Faculty Scholar in 2008, UNC WOWS Scholar 2009-2011, IEEE VGTC Virtual Reality Technical Achievement Award in 2010, and many best paper awards at international conferences. She is a Fellow of ACM, IEEE, and Eurographics.

Her research interests include computational robotics, haptics, physically-based modeling, virtual reality, sound rendering, and geometric computing. She has (co-)authored more than 300 refereed publications in these areas and co-edited/authored four books. She has served on hundreds of program committees of leading conferences and co-chaired dozens of international conferences and workshops. She is currently a member of Computing Research Association-Women (CRA-W) Board of Directors, Chair of IEEE Computer Society (CS) Fellows Committee, Chair of IEEE CS Computer Pioneer Award, and Chair of ACM SIGGRAPH Outstanding Doctoral Dissertation Award. She is a former member of IEEE CS Board of Governors, a former Editor-in-Chief of IEEE Transactions on Visualization and Computer Graphics (2011-2014), a former Chair of IEEE CS Transactions Operations Committee, and a member of several editorial boards. She also has served on several steering committees and advisory boards of international conferences, as well as government and industrial technical advisory committees.

Accelerated Computing Using FPGAs

Viktor K. Prasanna
University of Southern California
Charles Lee Powell Chair in Engineering at the University of Southern California,
ACM Fellow and IEEE Fellow

Abstract

FPGAs have matured over the years and are being used along with multi-core and emerging memory technologies to realize advanced platforms to accelerate a variety of applications. This talk will review the promise of reconfigurable computing leading up to current trends in accelerators. We will illustrate FPGA-based parallel architectures and algorithms for a variety of data analytics kernels in advanced networking, streaming graph processing and machine learning. While demonstrating algorithm-architecture co-design methodology to realize high performance accelerators for deep packet inspection, regular expression matching, packet classification, traffic classification, heavy hitter detection, etc., we demonstrate the role of modeling and algorithmic optimizations to develop highly efficient IP cores. We also show high throughput and energy efficient accelerator designs for a class of graph analytics and machine learning kernels. Our approach is based on high level abstractions of the architectures and design of efficient data structures, algorithms and mapping methodologies. We illustrate the performance improvements such accelerators offer and demonstrate the suitability of accelerators for these computations. We conclude by identifying opportunities and challenges in exploiting emerging heterogeneous architectures composed of multi-core processors, FPGAs, GPUs and coherent memory.

Biography

Viktor K. Prasanna (ceng.usc.edu/~prasanna) is Charles Lee Powell Chair in Engineering in the Ming Hsieh Department of Electrical Engineering and Professor of Computer Science at the University of Southern California. He is the director of the Center for Energy Informatics at USC and leads the FPGA (fpga.usc.edu) and Data Science Labs (dslab.usc.edu). His research interests include parallel and distributed systems including networked sensor systems, embedded systems, configurable architectures and high performance computing. He served as the Editor-in-Chief of the IEEE Transactions on Computers during 2003-06 and is currently the Editor-in-Chief of the Journal of Parallel and Distributed Computing. Prasanna was the founding Chair of the IEEE Computer Society Technical Committee on Parallel Processing. He is the Steering Co-chair of the IEEE International Parallel and Distributed Processing Symposium (www.ipdps.org) and the Steering Chair of the IEEE International Conference on High Performance Computing (www.hipc.org). He is a Fellow of the IEEE, the ACM and the American Association for Advancement of Science (AAAS). He is a recipient of 2009 Outstanding Engineering Alumnus Award from the Pennsylvania State University. He received the 2015 W. Wallace McDowell award from the IEEE Computer Society for his contributions to reconfigurable computing.

Supercomputing: Past, Present and Future

Marc Snir
University of Illinois at Urbana-Champaign
Michael Faiman Professor at the University of Illinois at Urbana-Champaign, ACM Fellow and IEEE Fellow

Abstract

For three decades, High Performance Computing has pursued one simple strategy: Platforms have been built as clusters of commodity systems, chosen for their superior cost/performance. Customization has been limited to packaging, interconnect and "glue" software. The improvements in the performance of HPC platforms have been due to the cost/performance improvements of commodity hardware predicted by Moore's Law, and to the increasing size and cost of leading supercomputers.

As Moore's Law comes to an end, the main driver to increased HPC performance is stalling. Continued performance improvements will require qualitatively different approaches. These include the use of highly heterogeneous architectures and specialized compute engines; and the aggressive use of energy saving technologies and packaging that enables high-energy densities. These changes in the underlying hardware will require significant changes in HPC software, in order to avoid the inefficiencies that occur at the multiple interfaces between software layers. Such changes may come at the expense of programmability and portability.
Supercomputers will bear less similarity to mainstream computers and will become unique high-end scientific instruments.

The talk will discuss the evidence for the thesis outlined in the previous paragraph; provide some examples of areas where significant software changes are likely to occur; and discuss some of the economic and political implications of a new HPC reality.

Biography

Marc Snir is Michael Faiman Professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign. He currently pursues research in parallel computing.

He was Director of the Mathematics and Computer Science Division at the Argonne National Laboratory from 2011 to 2016 and head of the Computer Science Department at Illinois from 2001 to 2007. Until 2001 he was a senior manager at the IBM T. J. Watson Research Center where he led the Scalable Parallel Systems research group that was responsible for major contributions to the IBM SP scalable parallel system and to the IBM Blue Gene system.

Marc Snir received a Ph.D. in Mathematics from the Hebrew University of Jerusalem in 1979, worked at NYU on the NYU Ultracomputer project in 1980-1982, and was at the Hebrew University of Jerusalem in 1982-1986, before joining IBM. Marc Snir was a major contributor to the design of the Message Passing Interface. He has published numerous papers and given many presentations on computational complexity, parallel algorithms, parallel architectures, interconnection networks, parallel languages and libraries and parallel programming environments.

Marc is AAAS Fellow, ACM Fellow and IEEE Fellow. He has Erdos number 2 and is a mathematical descendant of Jacques Salomon Hadamard. He recently won the IEEE Award for Excellence in Scalable Computing and the IEEE Seymour Cray Computer Engineering Award.

Decomposed Fuzzy Systems

Shun-Feng Su
National Taiwan University of Science and Technology
Chair Professor of Electrical Engineering at the National Taiwan University of Science and Technology, IEEE Fellow

Abstract

In the talk, a novel fuzzy structure termed as the decomposed fuzzy system (DFS) is proposed to act as the fuzzy approximator. The proposed structure is to decompose each fuzzy variable into layers of fuzzy systems and each layer is to characterize one traditional fuzzy set. Similar to forming fuzzy rules in traditional fuzzy systems, layers from different variables will form the so-called component fuzzy systems. The structure of DFS is proposed to facilitate minimum distribution learning effects among component fuzzy systems so that the learning can be very efficient. It can be seen from our experiments that even when the rule number increases, the learning time in terms of cycles is still almost constant. It can also be found that the function approximation capability and learning efficiency of the DFS are much better than that of the traditional fuzzy systems when employed in adaptive fuzzy control systems. Besides, in order to further reduce the computational burden, a simplified DFS is proposed in this study to satisfy possible real time constraints required in many applications. From our simulation results, it can be seen that the simplified DFS can perform fairly with a more concise decomposition structure. Furthermore, when used in modeling, the proposed DFS not only can have much faster convergent speed, but also can achieve a smaller testing error than those of other fuzzy systems.

Biography

Shun-Feng Su received the B.S. degree in electrical engineering, in 1983, from National Taiwan University, Taiwan, R.O.C., and the M.S. and Ph.D. degrees in electrical engineering, in 1989 and 1991, respectively, from Purdue University, West Lafayette, IN.

He is now a Chair Professor of the Department of Electrical Engineering, National Taiwan University of Science and Technology, Taiwan, R.O.C. He is an IEEE Fellow and CACS fellow. He has published more than 300 refereed journal and conference papers in the areas of robotics, intelligent control, fuzzy systems, neural networks, and non-derivative optimization. His current research interests include computational intelligence, machine learning, virtual reality simulation, intelligent transportation systems, smart home, robotics, and intelligent control.

Dr. Su is very active in various international/domestic professional societies. He is now the past president of the International Fuzzy Systems Association. He also serves as a board member of various academic societies. Dr. Su also acted as General Chair, Program Chair, or various positions for many international and domestic conferences. Dr. Su currently serves as Associate editors of IEEE Transactions on Cybernetics, IEEE Transactions on Fuzzy Systems, and IEEE Access, a subject editor (Electrical Engineering) of the Journal of the Chinese Institute of Engineers, and the Editor-in-Chief of International Journal of Fuzzy Systems.

Using Kernels to Harness the Complexity of Big Data Applications

Benjamin W. Wah
Chinese University of Hong Kong, China
Provost and Wei Lun Professor at the Chinese University of Hong Kong, ACM Fellow and IEEE Fellow

Abstract

Big Data is emerging as one of the hottest multi-disciplinary research fields in recent years. Big data innovations are transforming science, engineering, medicine, healthcare, finance, business, and ultimately society itself. In this presentation, we examine the key properties of big data (volume, velocity, variety, veracity, and value) and their relation to some applications in science and engineering. To truly handle big data, new paradigm shifts will be necessary. Successful applications in big data will require in situ methods to automatically extracting new knowledge from big data, without requiring the data to be centrally collected and maintained. Traditional theory on algorithmic complexity may no longer hold, since the scale of the data may be too large to be stored or accessed. To address the potential of big data in scientific discovery, challenges on data complexity, computational complexity, and system complexity will need to be solved. We propose a new approach based on identifying kernels to harness the complexity of big data applications. Kernel data is a compact and manageable representation of the original data, with similar structure, data properties, or meta-properties. We illustrate these challenges and approaches by drawing on examples in various applications in finance and multimedia.

Biography

Benjamin W. Wah is currently the Provost and Wei Lun Professor of Computer Science and Engineering of the Chinese University of Hong Kong, as well as the Chair of the Research Grants Council of the University Grants Committee, Hong Kong, and the Franklin W. Woeltge Emeritus Professor of Electrical and Computer Engineering, University of Illinois, Urbana-Champaign. Before then, he served as the Director of the Advanced Digital Sciences Center in Singapore, as well as the Franklin W. Woeltge Professor of Electrical and Computer Engineering and Professor of the Coordinated Science Laboratory of the University of Illinois, Urbana-Champaign, IL. He received his Ph.D. degree in computer science from the University of California, Berkeley, CA, in 1979. He has received numerous awards for his contributions, which include the IEEE CS Technical Achievement Award (1998), the IEEE Millennium Medal (2000), the IEEE-CS W. Wallace-McDowell Award (2006), the Pan Wen-Yuan Outstanding Research Award (2006), the IEEE-CS Richard E. Merwin Award (2007), the IEEE-CS Tsutomu Kanai Award (2009), and the Distinguished Alumni Award in Computer Science of the University of California, Berkeley (2011). Wah's current research interests are in the areas of big data applications and multimedia design and processing

Wah cofounded the *IEEE Transactions on Knowledge and Data Engineering* in 1988 and served as its Editor-in-Chief between 1993 and 1996. He currently serves as the Honorary Editor-in-Chief of *Knowledge and Information Systems* and is on the editorial boards of *Information Sciences*, *International Journal on Artificial Intelligence Tools*, *Journal of VLSI Signal Processing*, *World Wide Web,* and *Journal of Computer Science and Technology*. He has served the IEEE Computer Society in various capacities, including Vice President for Publications (1998 and 1999) and President (2001). He is a Fellow of the AAAS, ACM, and IEEE.

Several Thoughts About Network Computing Paradigms

Yaoxue Zhang
Central South University
Former President of Central South University, Academician of Chinese National Academy of Engineering

Abstract
Network computing and its service model have experienced rapid development during the past decade. However, as the era of Internet-of-Things (IoT) is approaching, current network computing paradigms and service models are facing significant challenges for supporting lightweight IoT devices. In this talk, I will introduce a new service model based on transparent computing, named Block-stream as a Service (BaaS), for IoT scenarios, and then describe our preliminary implementations for some typical IoT devices as well. Detailly, for lightweight IoT devices, we propose BOAT, a block-streaming machine code execution method for ELF-based Apps based on transparent computing, which can efficiently achieve flexible service provisioning with improved scalability and reduced service loading delay. As for Android devices, we propose a BaaS technique for fast App deployment based on Android compilation mechanism, which reduces the time cost of App installation so that provides more efficient on-demand App services for Android users. Finally, I will discuss some future developments and collaborations of transparent computing, cloud computing, and Big Data.

Biography
Yaoxue Zhang received his B.S. degree from Northwest Institute of Telecommunication Engineering, China, in 1982, and his Ph.D. degree in computer networking from Tohoku University, Japan, in 1989. Currently, he is a professor in the School of Information Science and Engineering at Central South University, China, and also a professor in the Department of Computer Science and Technology at Tsinghua University, China. His research interests include computer networking, operating systems, ubiquitous/ pervasive computing, transparent computing, and big data. He has published over 200 technical papers in international journals and conferences, as well as 9 monographs and textbooks. He is a fellow of the Chinese Academy of Engineering and the EiC of Chinese Journal of Electronics.

计算机科学最新进展
高端研讨会

High-end Forum on Latest Development in Computer Science

June 19 2018

Conference Hall

Second Floor

National Supercomputing Center in Changsha

主办：“111计划”高效能计算学科创新引智基地
湖南大学信息科学与工程学院
国家超级计算长沙中心