

International Workshop ²⁰¹⁹ on Advanced Infrastructure Software & Distributed Computing

基础软件与分布式计算先进技术研讨会2019

December 17~ 18, 2019
in Shanghai



THE UNIVERSITY
of EDINBURGH



上海交通大学
SHANGHAI JIAO TONG UNIVERSITY



HUAWEI

Registration 9:00 AM - 9:25 AM

Opening Remarks 9:25 AM - 9:30 AM

Welcome

Morning talks: (Moderator: Sihao Li, Huawei) 9:30 AM - 12:10 AM

[Analyses for Social Media Marketing ---By Wenfei Fan, University of Edinburgh](#)

[Building Low-latency and Scalable Systems Using Non-volatile Memory---By Haibo Chen, Shanghai Jiao Tong University & Huawei](#)

Poster Session & Break 10:50 AM - 11:30 AM

[Data-Centric Parallelization---By Björn Franke, University of Edinburgh](#)

Lunch 12:10 AM – 1:30 PM

Afternoon talks: (Moderator: TBD, Shanghai Jiaotong University) 1:30 PM - :20 PM

[Business and Challenges of Compilers and Programming Languages Lab---Xinyu Feng, Nanjing University & Huawei](#)

[TSTL: a novel logic for analysing the spatio-temporal properties of spatial stochastic systems---Jane Hillston, University of Edinburgh](#)

Poster Session & Break 2:50 PM – 3:20 PM

[Knowledge Base Refinement and Enhancement---By Lei Chen, Hong Kong University of Science and Technology & Huawei](#)

[Semantic Parsing and Dialog for Knowledge-Graph Query---By Mark Steedman, University of Edinburgh](#)

[Knowledge Graph: Where are we going from here?---By Jeff Pan, University of Aberdeen & Huawei](#)

Closing Remarks 5:20 PM – 5:30 PM

Banquet 6:00 PM (Invited only)

Registration 9:00 AM - 9:30 AM

Session Discussion 9:30 – 12:30

<p>Data Management (Moderator: Lei Chen & Sihao Li) Room: 101</p>	<p>Compiler and CPL (Moderator: Xinyu Feng & Yaoqing Gao) Room: 103</p>	<p>Distributed Computing & AI (Moderator: Fan Yu) Room: 102 (大会议室)</p>
<p>Huawei: 周敏奇(Minqi Zhou)、罗正海(Zhenghai Luo)、张琼(Qiong Zhang); 朱建峰(Jianfeng Zhu)、曹琛(Chen Cao)、王云鹏(Yunpeng Wang)</p>	<p>Huawei: 王学智(Xuezhi Wang)、海丽娟(Lijuan Hai)、杨扬(Yang Yang)、魏伟(Wei Wei)、何璐(Lu He)、曾建江(Jianjiang Zeng)</p>	<p>Huawei: Yantao Jia, Jeff Pan, 文庆(Qing Wen), 周杰(Jie Zhou), <u>TBD</u></p>
<p>UofE: Wenfei Fan, Peter Buneman, Paolo Guagliardo,</p>	<p>UofE: Björn Franke, James Cheney, Hugh Leather</p>	<p>UofE: Björn Franke, Shay Cohen, Matthieu Labeau, Kai Xu, Mahesh Marina, Amos Storkey, Pramod Bhatotia, Hakan Bilen, Tiejun Ma</p>
<p>SJTU: 如果您有兴趣参加请反馈给联系人 (最后一页)</p>	<p>SJTU:</p>	<p>SJTU:</p>

Poster Session & Break 10:40 AM - 11:20 AM

Steering Committee Meeting
Room: 107

Lunch 12:10 AM – 1:30 PM (Arranged separately by each session)



Analyses for Social Media Marketing

Abstract:

Social media marketing promotes products or services by capitalizing on associations among entities in social networks. It has proven effective, and even fatal, in commercial, online public sentiment and politics; it is considered a killer application of big data analytics. No matter how important, the study of associations is challenging. Is it possible to unify logic rules and machine learning to accurately capture associations? How can we efficiently discover useful associations from a social network with billions of people and trillions of links? How can we leverage associations to identify potential customers in large-scale social networks? Moreover, can we still practice social media marketing when we have no access to the data of social networks? This talk tackles these issues.



Wenfei Fan
University of Edinburgh

Bio:

Wenfei Fan is Professor of Web Data Management at the University of Edinburgh. His research investigates database theory and database systems. He has both formalised the problems of querying big data and has developed radically new techniques that overcome the limits associated with conventional database systems. His work has been adopted in the telecommunications industry for analyzing massive data sets that defied their current technology. In addition, Fan has made seminal contributions to data quality, in which he devised new techniques for data cleaning that have found wide commercial adoption. He has also contributed to our understanding of semi-structured data. Fan is a foreign member of the Chinese Academy of Sciences, a fellow of the Royal Society, a fellow of the Academia Europæa, a Fellow of the Association for Computing Machinery, a Fellow of the Royal Society of Edinburgh and winner of the Roger Needham Award in 2008.

Building Low-latency and Scalable Systems Using NVM

Abstract:

The recent availability of Intel/Micron's 3D Xpoint to the market marks the transition of non-volatile memory (NVM) from research artifact to real product. While NVM shares similarities like persistency and byte-addressability with DRAM, it has a different performance characteristic with DRAM and changes the persistency boundary of computer memory hierarchy. In this talk, I will first present a characterization on the performance of NVM and then presents a series of work that exploits NVM to build low-latency and scalable computing systems, including file systems, Java virtual machines and transaction-processing systems. Finally, I will discuss the challenges and opportunities ahead and outline some future directions..



Haibo Chen,
Shanghai
Jiao Tong University
& Huawei

Bio:

Haibo Chen is a Professor at School of Software, Shanghai Jiao Tong University, where he co-founds and leads the Institute of Parallel and Distributed Systems (IPADS) . He is currently on sabbatical at Huawei, where he founds and directs the OS Kernel Lab and serves as the Chief Scientist for OS. Haibo's main research interests are building scalable and dependable systems software, by leveraging cross-layering approaches spanning computer hardware, system virtualization and operating systems. His research work has been widely adopted by open-source community like Linux and Open JDK as well as commercial products deployed over hundreds of billions of devices, resulted in numerous publications in top systems conference like SOSP/OSDI. He currently chairs ACM SIGOPS ChinaSys, continuously serves on the program committees of SOSP/CCS/IEEE S&P/ASPLOS/ISCA/PPoPP/Usenix ATC/EuroSys as well as the editorial boards of ACM Transactions on Storage and Communications of the ACM. He is a Distinguished Member of ACM and CCF.

Data-Centric Parallelization

Abstract:

We take a fresh look at automatic parallelization techniques and identify a number of conceptual problems, which limit the success of traditional compiler-based parallelization approaches. In the first part of our presentation we show how dynamic information, e.g. obtained from memory access profiling, can complement static analyses, while machine learning techniques outperform manually developed code mapping techniques. In the second part we focus on novel models of parallelism and parallelization analyses based on commutativity and parallel algorithmic skeletons. We show that this high-level approach leads to the detection of more parallelism than what is possible with traditional dependence based techniques. We round off our talk with an outlook at combined source code rejuvenation and data abstraction techniques, and how these can be used to elevate parallelizing compilers to the next level.



Björn Franke,
University of Edinburgh

Bio:

Björn Franke is a Reader in the School of Informatics. His research interests include software transformation driven by dynamic information. He has published extensively on iterative and machine learning based compiler optimisation, hardware accelerated dynamic binary translation and profile-driven parallelization.

Semantic Parsing and Dialog for Knowledge-Graph Query

Abstract:

The talk discusses a variety of techniques for using semantic parsing to interrogate existing and envisioned databases and knowledge graphs in natural language, and in particular its use for open domain-question answering from text at web scale, contrasting deep learning and neural network computation with parser-based machine-reading methods.

As soon as question answering of any of these kinds is supported, users will expect to be able to ask follow-up questions, and to elaborate their questions in an interactive dialog with the system. Such dialogs raise some very old problems of co-reference and user intention recognition. Deep learning and neural computation offer some new and promising techniques for problems such as the alignment of antecedents and co-referring pronouns.

All such dialogs are at an abstract level negotiations concerning the partition of resources, notably the participant's time and computational effort. Such negotiation dialog skills are notoriously hard to acquire by machine learning, with supervised learning prone to overfit, and reinforcement by results largely ineffective. Recent models that have been very successful in game playing programs like DeepMind AlphaGo combine a supervised policy for proposing moves with "roll-out" of expected outcomes using hand-built or reinforcement rules and Minimax search for evaluating those moves. This architecture can be applied to negotiation dialogues, including those resulting from conversational open-domain QA using semantic parsing. If I have time, I will conclude by considering some possibilities for extending these dialog capability to speech assistants on mobile devices.



Mark Steedman,
University of Edinburgh

Bio:

Steedman currently holds the Chair of Cognitive Science in the School of Informatics at the University of Edinburgh. He works in computational linguistics, artificial intelligence, and cognitive science, on the Generation of Meaningful Intonation for Speech by Artificial Agents, Animated Conversation, The Communicative Use of Gesture, Tense and Aspect, and Combinatory Categorical Grammar (CCG). He is also interested in Computational Musical Analysis and Combinatory Logic. Steedman is a Member of the Academia Europæa, a Fellow of the British Academy, a Fellow of the Royal Society of Edinburgh and an AAI Fellow. In 2018 he was the recipient of the Lifetime Achievement Award from the Association for Computational Linguistics.

TSTL: a novel logic for analysing the spatio-temporal properties of spatial stochastic systems

Abstract:

We consider systems in which dynamic behaviour is spatially distributed and also stochastic. In particular the agents within the system have a location, and the ability of agents to interact may depend on agents being co-located or within a given range of each other. Such systems are typically studied by simulations. Current simulation-based approaches provide summary information about the satisfaction of properties over the spatial domain, providing estimated values that include intrinsic uncertainty. In this work we seek to take this uncertainty into account and to enrich the summary information through the use of a novel logic, the Three-Valued Spatio-Temporal Logic (TSTL), which allows us to reason, not only about the behaviour of the system, but also about the evolution of the satisfaction of properties expressed in a spatio-temporal logic. This provides additional insight into the dynamic behaviour of the system under study. For example, in the analysis of the efficacy of a control measure for fire spread, we can verify whether the spread in a specific area will happen with probability under a given threshold over time. We can also identify the locations that are at highest risk, because they are surrounded by locations with high probability of burning. We use a three-valued logic, with the third value unknown indicating when there is insufficient evidence to make a judgement about a property. This can be taken as an indication of when more simulation trajectories are needed to evaluate propositions more precisely.

(Joint work with Ludovica Luisa Vissat, Michele Loreti, Laura Nenzi and Glenn Marion).



Jane Hillston,
University of Edinburgh


Bio:

Jane Hillston is Professor of Quantitative Modelling at the University of Edinburgh and is currently Head of the School of Informatics at the University of Edinburgh. Her work on the stochastic process algebra PEPA was recognised by the British Computer Society in 2004, which awarded her the first Roger Needham Award. She was elected to fellowship of the Royal Society of Edinburgh in 2007 and she is a member of Academia Europaea. Her principal research interests are in formal approaches to modelling the dynamic behaviour of systems, particularly performance modelling. Recently this has included making fluid approximations of discrete state systems to tackle the problems in state space explosion. She has also been working on the integration of inference into formal modelling based on process algebra.

Knowledge Base Refinement and Enhancement

Abstract:

Nowadays knowledge bases has been serving as promising and effective tools for people to explore knowledge in different areas, such as DBPedia, Yago, Freebase and Knowledge Vault from Google. Therefore, enhancing the knowledge bases and making a good trade-off between completeness and correctness has been raised up as an important task. In order to further increase the utility of knowledge bases, various refinement methods have been proposed. These methods either try to add or infer missing knowledge to the original knowledge bases or aim at identifying incorrect information contained. In this talk, I will provide an overview of the knowledge base refinement and enhancement approaches. Then, I will discuss knowledge base enhancement via data facts and crowdsourcing, knowledge truth discovery from conflicting sources and knowledge bases canonicalization from source texts. Finally, I will present some interesting work on subjective knowledge base construction and enhancement.



Lei Chen,
Hong Kong University
of Science and
Technology & Huawei


Bio:

Lei Chen received the BS degree in computer science and engineering from Tianjin University, Tianjin, China, in 1994, the MA degree from Asian Institute of Technology, Bangkok, Thailand, in 1997, and the PhD degree in computer science from the University of Waterloo, Canada, in 2005. He is currently a full professor in the Department of Computer Science and Engineering, Hong Kong University of Science and Technology. His research interests include data driven machine learning, knowledgebase construction and refinement, crowdsourcing over social media, social media analysis, probabilistic and uncertain databases, and privacy preserved data publishing. The system developed by his team won the excellent demonstration award in VLDB 2014. He got the SIGMOD Test-of-Time Award in 2015. He is PC Co-Chair for VLDB2019, PC track chairs for SIGMOD 2014, VLDB 2014, ICDE 2012, CIKM 2012, SIGMM 2011. He has served as PC members for SIGMOD, VLDB, ICDE, SIGMM, and WWW. Currently, he serves as Editor-in-Chief of VLDB Journal and an associate editor-in-chief of IEEE Transaction on Data and Knowledge Engineering. He is an executive member of the VLDB endowment, Fellow of IEEE and ACM Distinguished Scientist.

Knowledge Base Refinement and Enhancement

Abstract:

The mission of Compilers and Programming Languages Lab (the Lab thereafter) at Central Software Institute, Huawei 2012 Lab is to deliver extreme performance on Huawei AKKD chips series, and to provide high productivity for developers. The lab has been actively developing new compilers, language virtual machines and programming languages to support Huawei consumer devices, ICT businesses, Huawei cloud, etc. In this talk I'll give an overview of the past and ongoing efforts in the Lab. Then I'll introduce the key challenges and future collaboration opportunities, in particular in the areas of AI programming, probabilistic programming, heterogeneous and parallel programming.



Xinyu Feng,
Nanjing University
& Huawei

Bio:

Xinyu Feng is the director of the Programming Languages Lab at Central Software Institute, Huawei 2012 Lab. He is also a professor at Nanjing University. He got his PhD from Yale University in 2007, and worked at Toyota Technological Institute at Chicago, and University of Science and Technology of China (USTC) before joining Nanjing University in 2018. His research interests are in the areas of programming language semantics and formal program verification..

Knowledge Graph: Where are we going from here?

Abstract:

Knowledge Graph has become popular in data representation and knowledge management applications widely across search engine, finance, healthcare, Open Science and Industrial 4.0 in recent years. It is regarded as a backbone for AI systems, providing an umbrella overlay for cross-walks across siloed data resources as well as a semantic bridge between AI systems and human users. In this talk, I will firstly introduce the notion of Knowledge Graph (What is it), and then discuss some typical application scenarios of Knowledge Graph (What problems does it solve) and related research challenges. I will conclude the talk by sharing some thoughts on where Knowledge Graph might lead us into the near future.

Bio:



Jeff Pan,
University of
Aberdeen & Huawei

Dr Jeff Z. Pan received his Ph.D. in computer science from The University of Manchester in 2004. He is a Reader in the Department of Computing Science at University of Aberdeen and the Director of the Aberdeen-Wuhan Joint Research Lab on Knowledge Engineering and Information Security. He was an official reviewer of the international Knowledge Graph standards RDF and SPARQL and was a key contributor of the international standard (OWL) of Knowledge Graph schemas. He leads the development of the award-winning TrOWL approximate reasoner, which is one of the top three OWL 2 DL reasoners in the sound and complete Ontology Reasoner Evaluation (ORE2014). He was the Chief Scientist of the EU Marie-Curie K-Drive project. His research focuses primarily on knowledge representation and artificial intelligence, in particular on knowledge graph based reasoning and learning, as well as their applications. He is currently on sabbaticals in Huawei's Edinburgh Research Centre. He is an Associate Editor of the Journal of Web Semantics (JWS) and of the International Journal on Semantic Web and Information Systems (IJSWIS). He is a PC Co-Chair ISWC2020, as well as RR2007 and JIST2011. He was a General Co-Chair of JIST2014, CCKS2014 and RW2016. He is the Chief Editor of the first two books on Knowledge Graph.



Please email to liuxiaoli@huawei.com or call +86 13510712796 if any questions, thanks

有任何问题都可以联系 刘晓莉 00238413 , 邮箱 liuxiaoli@huawei.com , 电话 : 13510712796 , 谢谢。

Thank you.

把数字世界带入每个人、每个家庭、
每个组织，构建万物互联的智能世界。

Bring digital to every person, home, and
organization for a fully connected,
intelligent world.

**Copyright©2018 Huawei Technologies Co., Ltd.
All Rights Reserved.**

The information in this document may contain predictive statements including, without limitation, statements regarding the future financial and operating results, future product portfolio, new technology, etc. There are a number of factors that could cause actual results and developments to differ materially from those expressed or implied in the predictive statements. Therefore, such information is provided for reference purpose only and constitutes neither an offer nor an acceptance. Huawei may change the information at any time without notice.

