

CS Seminar

Exploring the intersection of robust system design and machine learning

Dr Yanjing Li
Department of Computer Science
University of Chicago

Date:
July 3, 2018
Tuesday
2:00 pm

Venue:
Room 328
Chow Yei Ching Building
The University of Hong Kong

Abstract:

Rapid advances in computing systems have transformed every aspect of life as we know it. However, we are now facing significant challenges in our quest for even more advanced computing systems. With significant vulnerability to failures and defects in CMOS and emerging technologies, hardware robustness is a key challenge for a large class of future computing systems — from edge devices all the way to cloud servers. Due to explosive growth in our dependency and demands on these systems, there is an urgent need to design robust systems that performs correctly despite underlying disturbances caused by hardware failures, design flaws, software bugs, environmental effects, and malicious attacks.

At the same time, exciting opportunities in robust system design also arise with innovations in new technologies and applications. In particular, machine learning has already achieved substantial breakthroughs in many computing domains and is expected to become even more prominent in the future. In this talk, we will explore the intersection of robust system design and machine learning from two different angles.

First, existing machine learning techniques may be effectively utilized to design efficient and low-cost robust systems. We will show an example where machine learning is used to guide dynamic soft error resilience tuning in microprocessors, leading to 2X improvement in overall energy efficiency compared to static hardening techniques, which up to now have been shown to be the one of the most efficient and effective soft error resilience approaches, without sacrificing reliability.

Second, robust systems optimized for efficient processing of machine learning applications are critical for pushing the frontiers of these applications. We will discuss our work on a direct-modulated optical interconnection network for large-scale interposer systems. Using multi-chip module GPUs as a case study, we find that our network design is capable of scaling up the number of streaming multiprocessors by up to 64X compared to the state-of-the-art today, while outperforming various competing designs in terms of energy efficiency, performance, and reliability. This will help satisfy the computing demands from future machine learning and other emerging applications.

About the Speaker:

Dr. Yanjing Li is an Assistant Professor in the Department of Computer Science (Systems Group) at the University of Chicago. Prior to joining University of Chicago, she was a senior research scientist at Intel Labs. She received a Ph.D. in Electrical Engineering from Stanford University, and a M.S. in Mathematical Sciences (with honors) and a B.S. in Electrical and Computer Engineering (with a double major in Computer Science) both from Carnegie Mellon University. Her research interests lie broadly in computer architecture, emerging technologies, and VLSI design and validation. The focuses of her current research include interactions between computing systems and machine learning, photonic interconnects and processing, hardware security, and robust memory systems. She has won various awards including the NSF/SRC Energy-Efficient Computing: from Devices to Architectures (E2CDA) program, Intel Labs Gordy Academy Award (highest honor in Intel Labs), multiple Intel recognition awards, Outstanding Dissertation Award (European Design and Automation Association), Best Student Paper Award (IEEE International Test Conference), and the Best Paper Award (IEEE VLSI Test Symposium).

All are welcome!

For enquiries, please call 2859 2180 or email
enquiry@cs.hku.hk
Department of Computer Science
The University of Hong Kong

