COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF MICHIGAN

DISSERTATION DEFENSE

Haojie Ye

Combating Randomness: Towards Efficient Data-Intensive Applications Using Software-Hardware Co-Design

Tuesday, September 3, 2024
11:00am – 1:00pm
4901 Beyster
Hybrid – Zoom

ABSTRACT: The emergence of data-intensive applications has a tremendous impact on our daily life. For example, social media processes billions of mutual connections and intelligently identifies potential relations. Web vendors analyze millions of commodities such as news, products, and services to produce personalized recommendations for users. Financial and medical institutions launch millions of financial and medical record transactions in secure cloud databases per day. However, all of the above use cases involve a large amount of data accesses that are considered irregular or even random.

Traditional computing platform architectures (e.g., CPUs and GPUs) have been shown to be inefficient for processing these data-intensive applications because (a) the large and irregular memory footprint renders the on-chip cache ineffective, (b) segmented and irregular cache line granularity accesses do not saturate the main memory bandwidth, and (c) the dataflow can be data dependent with frequent processor stalls.
To advance data-intensive applications, there is a critical need to synergize new ideas in hardware, algorithms, and systems to efficiently execute these vital applications. This thesis proposes a series of works that optimizes the performance of executing a wide range of emerging data-intensive applications. Specifically, we combat the irregularity and randomness in these applications using different software-hardware co-designs.

CHAIR: Prof. Trevor Mudge and Dr. Nishil Talati