COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF MICHIGAN

DISSERTATION DEFENSE

Armand Behroozi

Data-Aware Computing for Efficient Extended Reality Systems

Friday, January 17, 2025
3:00pm – 5:00pm
3725 Beyster
Hybrid – Zoom Passcode: 686767

ABSTRACT: Augmented and virtual reality (AR/VR) or extended reality (XR) devices have the potential to fundamentally transform how people interact with the digital and physical worlds. These technologies aim to create fully immersive virtual environments and seamlessly integrate digital content into the physical world, enabling applications in education, healthcare, remote work, and design. However, a major bottleneck to the widespread use of these devices is their strict power and weight constraints. In order to create a compelling user experience, XR devices must be comfortable to wear for extended periods of time without overheating while simultaneously running many computationally expensive tasks.

The diminishing returns of Moore's Law and Dennard Scaling further compound these challenges. While specialized hardware accelerators offer potential efficiency gains, the rapidly evolving nature of XR workloads makes such solutions inflexible and prone to obsolescence.
Similarly, offloading computation to external resources introduces privacy concerns and depends on reliable, high-bandwidth wireless networks, which may be unavailable in environments such as urban canyons. For augmented reality, 100x improvements in power efficiency are necessary to support fully on-device computation. These challenges underscore the need for alternative approaches to optimize on-board processing.

This thesis presents application-specific software techniques to improve the efficiency of XR systems. First, we introduce Loner, an auto-vectorization technique that optimizes general-purpose compute workloads in XR devices by exploiting underutilized single instruction multiple data CPU resources. Second, we propose SlimSLAM, a domain-specific runtime scheduler for simultaneous localization and mapping (SLAM), used for head tracking in XR headsets. SlimSLAM reduces unnecessary computation by dynamically adjusting over-provisioned algorithmic parameters based on real-time state information. Finally, we present SlimDepth, an intelligent keyframe selection method for depth estimation, enabling virtual objects to interact with the physical world while minimizing computation.

CHAIR: Prof. Scott Mahlke