



University Seminar Series: GW Robotics, Artificial Intelligence (AI), and Transportation for Smart Health & Smart Cities (GRITS)

Linearized Tensor Format for Performance-Portable Sparse Tensor Computation

Dr. Jeewhan Choi

**Assistant Professor, Electrical Engineering
Oregon University**

Many critical application domains, such as healthcare, cybersecurity, and data mining, produce and manipulate massive amounts of high-dimensional data. Such datasets can naturally be represented as sparse tensors. A fundamental challenge in sparse tensor computation involves determining how to organize the data, as to minimize storage, improve data locality, increase parallelism, and decrease workload imbalance and synchronization. To address this challenge, we present a novel sparse tensor storage format that achieves these goals to deliver an order of magnitude faster execution time for a key sparse tensor operation compared to prior state-of-the-art formats. Furthermore, we demonstrate that our new format operates efficiently across vastly different parallel computing architectures with minimal adaption.



Prof. Jeewhan Choi is an Assistant Professor in the Department of Computer Science at the University of Oregon. During his PhD, he worked on designing parallel and scalable algorithms for scientific applications and modeling their performance and energy efficiency on the latest high-performance computing (HPC) systems. After graduation, he worked as a research staff member at the IBM T. J. Watson Research Center on designing and optimizing tensor decomposition algorithms for Big Data analytics.

Wednesday, May 1, 2024

2-3pm

SEH 2000