

*Distinguished Speakers Series Presents*

**Xiaodong Zhang, Ohio State University**



**Fast Accesses in Memory and Storage in Big Data Environment**

A major goal of algorithms analysis and implementation in data processing is to read and write data records from memory or storage in high speed at a low cost for a given data storage format. As the data volume generated in the society continues to grow in an increasingly rapid way, we have reevaluated several commonly used data accessing methods including LSM-tree for sequentially archived data, and storing/retrieving methods for key-value stored data. In this talk, I will show their limits and inabilities to handle big volume of data in a scalable way. I will also present three new research results: (1) re-enabling buffer caching capability for LSM-tree to achieve high performance of both reads and writes to process sequentially archived data, (2) balancing both network bandwidths and storage transfers for relational tables in large clusters, and (3) maximizing throughput of in-memory key-value stores by GPUs. All the related algorithms and software implementations are open sourced, some of which have been adopted in production systems.

Bio: Xiaodong Zhang is the Robert M. Critchfield Professor in Engineering and Chair of the Computer Science and Engineering Department at the Ohio State University. His research interests focus on data management in computer and distributed systems. He has made strong efforts to transfer his academic research into advanced technology to update the design and implementation of major general-purpose computing systems. He received his Ph.D. in Computer Science from University of Colorado at Boulder, where he received Distinguished Engineering Alumni Award in 2011. He is a Fellow of the ACM, and a Fellow of the IEEE.

**Thursday, November 19, 2015**

**3:30–4:30 PM**

**University at Buffalo - North Campus - Knox 14**