# System, Architectural and Application level analysis of Big Data Applications for Performance and Energy-Efficiency

ECE PhD Seminar

Ph.D. Advisor:
Dr. Houman Homayoun

Monday, March 19, 2018
11:00 am-12:00 pm
ENGR 2901


## Abstract

Datacenters provide high performance and flexibility for users and cost efficiency for operators. Hyperscale datacenters are harnessing massively scalable computer resources for large-scale data analysis. However, cloud/datacenter infrastructure does not scale as fast as the input data volume and computational requirements of big data and analytics technologies. Thus, more applications need to share CPU at the node level that could have large impact on performance and operational cost. To address this challenge, this dissertation shows how concurrently fine-tune parameters at the application, microarchitecture, and system levels create opportunities to co-locate applications at the node level and improve energy-efficiency of the server while maintaining performance. Co-locating and self-tuning of unknown applications are challenging problems, especially when co-locating multiple big data applications concurrently with many tuning knobs, potentially requiring exhaustive brute-force search to find the right settings. This research challenge upsurges an imminent need to develop a technique that co-locates applications at a node level and predict the optimal system, architecture and application level configure parameters to achieve the maximum energy efficiency. Towards this goal, we develop an Energy-Efficient Co-Locating and Self-Tuning (ECoST) technique for co-located MapReduce applications to enhance their energy-efficiency.