



DISSERTATION DEFENSE

Nishil Talati



Optimizing Emerging Graph Applications Using Hardware-Software Co-Design

Tuesday, September 27, 2022

10:00 – 12:00pm

Hybrid– LEC4231 [Zoom](#)

Passcode: defense

ABSTRACT: A graph is a ubiquitous data structure that models entities and their interactions through the collections of nodes and edges. It is widely employed in several important application domains ranging from social media, navigation tools, search engines, physics simulations, and biology. Despite its prevalence, the performance of graph algorithms on commercial platforms is limited. This is mainly due to the irregular nature of memory accesses and convoluted control flow used in graph algorithms while accessing large amounts of real-world graph data (i.e., billions of nodes/edges). Therefore, there is a pressing need for optimizing the performance of graph algorithms.

In this thesis, I present a systematic optimization study of a variety of graph algorithms running on both static and dynamic graphs. At a high level, I first analyze the unique challenges and execution bottlenecks of the state-of-the-art graph software frameworks running on commercial hardware platforms. I then use the insights obtained from this analysis to propose design optimizations catered to the unique workload characteristics of a diversity of graph workloads.

Specifically, first, I propose Prodigy—a hardware-software co-design solution to improve the performance of traditional graph processing algorithms (e.g., PageRank and SSSP) on multi-core CPUs. Second, I present an in-depth study of random walk-based graph learning algorithms on temporal graphs (a type of dynamic graph). Specifically, this study delivers high-performance, open-source CPU and GPU implementations of important graph learning applications, conducts a detailed performance analysis, and makes recommendations for future optimizations. Third, I showcase NDMiner—a domain-specialized Near Data Processing (NDP) architecture that improves the performance of Graph Pattern Mining (GPM) workloads. Lastly, I present Mint—a novel hardware accelerator architecture and an accompanying programming model for efficiently mining motifs in temporal graphs.

CHAIR: Prof. Ronald Dreslinski